

# UNKs Everywhere: Adapting Multilingual Language Models to New Scripts

Jonas Pfeiffer<sup>1</sup>, Ivan Vulić<sup>2</sup>, Iryna Gurevych<sup>1</sup>, Sebastian Ruder<sup>3</sup>

<sup>1</sup>Ubiquitous Knowledge Processing Lab, Technical University of Darmstadt

<sup>2</sup>Language Technology Lab, University of Cambridge

<sup>3</sup>DeepMind

pfeiffer@ukp.tu-darmstadt.de

## Abstract

Massively multilingual language models such as multilingual BERT offer state-of-the-art cross-lingual transfer performance on a range of NLP tasks. However, due to limited capacity and large differences in pretraining data sizes, there is a profound performance gap between resource-rich and resource-poor target languages. The ultimate challenge is dealing with under-resourced languages not covered at all by the models and written in scripts *unseen* during pretraining. In this work, we propose a series of novel data-efficient methods that enable quick and effective adaptation of pretrained multilingual models to such low-resource languages and unseen scripts. Relying on matrix factorization, our methods capitalize on the existing latent knowledge about multiple languages already available in the pretrained model’s embedding matrix. Furthermore, we show that learning of the new dedicated embedding matrix in the target language can be improved by leveraging a small number of vocabulary items (i.e., the so-called *lexically overlapping* tokens) shared between mBERT’s and target language vocabulary. Our adaptation techniques offer substantial performance gains for languages with unseen scripts. We also demonstrate that they can yield improvements for low-resource languages written in scripts covered by the pretrained model.

## 1 Introduction

Massively multilingual language models pretrained on large multilingual data, such as multilingual BERT (mBERT; Devlin et al., 2019) and XLM-R (Conneau et al., 2020a) are the current state-of-the-art vehicle for effective cross-lingual transfer (Hu et al., 2020). However, while they exhibit strong transfer performance between resource-rich and similar languages (Conneau et al., 2020a; Artetxe et al., 2020), these models struggle with transfer to low-resource languages (Wu and Dredze, 2020) and languages not represented at all in their pre-

Amharic	Tibetan	Divehi
የአግዚአብሔርን	ལྷ་མཚན་ལྷ་མཚན་	މުޅިން ފަށައިގެން
የሚሰጠውን	བློ་བྲོ་	މި ދިވެހި ރާއްޖޭގެ
የአግዚአብሔር	བློ་བྲོ་	މި ދިވެހި ރާއްޖޭގެ
ቤተክርስቲያን	ལྷ་མཚན་	މި ދިވެހި ރާއްޖޭގެ
ኢትዮጵያውያን	ལྷ་མཚན་	މި ދިވެހި ރާއްޖޭގެ

Figure 1: Example tokens of unseen scripts.

training corpora (Pfeiffer et al., 2020b; Müller et al., 2021; Ansell et al., 2021). The most extreme challenge is dealing with *unseen languages with unseen scripts* (i.e., the scripts are not represented in the pretraining data; see Figure 1), where the pretrained models are bound to fail entirely if they are used off-the-shelf without any *further model adaptation*.

Existing work focuses on the embedding layer and learns either a new embedding matrix for the target language (Artetxe et al., 2020) or adds new tokens to the pretrained vocabulary. While the former has only been applied to high-resource languages, the latter approaches have been limited to languages with seen scripts (Chau et al., 2020; Müller et al., 2021) and large pretraining corpora (Wang et al., 2020). Another line of work adapts the embedding layer as well as other layers of the model via adapters (Pfeiffer et al., 2020b; Üstün et al., 2020). Such methods, however, cannot be directly applied to languages with unseen scripts.

In this work, we first empirically verify that the original tokenizer and the original embedding layer of a pretrained multilingual model fail for languages with unseen script. This implies that dedicated in-language tokenizers and embeddings are a crucial requirement for any successful model adaptation. The key challenge is aligning new target language embeddings to the pretrained model’s representations while leveraging knowledge encoded in the existing embedding matrix. We systematize existing approaches based on the pretrained information they utilize and identify *lexically overlapping tokens* that are present in both vocabularies as key carriers of such information (Søgaard et al.,

2018).<sup>1</sup> We then present novel, effective, and data-efficient methods for adapting pretrained multilingual language models to resource-low languages written in different scripts. Beyond lexical overlap, our methods rely on factorized information from the embedding matrix and token groupings.

We evaluate our approaches in the named entity recognition (NER) task on the standard WikiAnn dataset (Rahimi et al., 2019) and Dependency Parsing (DP; Nivre et al., 2016). We use 4 diverse resource-rich languages as source languages, and transfer to 17 and 6 resource-poor target languages respectively, including 5 languages with unseen scripts (Amharic, Tibetan, Khmer, Divehi, Sinhala). We show that our adaptation techniques offer unmatched performance for languages with unseen scripts. They also yield improvements for low-resource and under-represented languages written in scripts covered by the pretrained model.

**Contributions.** **1)** We systematize and compare current model adaptation strategies for low-resource languages with seen and unseen scripts. **2)** We measure the impact of initialization when learning new embedding layers, and demonstrate that non-random initialization starting from a subset of seen lexical items (i.e., *lexically overlapping* vocabulary items) has a strong positive impact on task performance for resource-poor languages. **3)** We propose methods for learning low-dimensional embeddings, which reduce the number of trainable parameters and yield more efficient model adaptation. Our approach, based on matrix factorization and language clusters, extracts relevant information from the pretrained embedding matrix. **4)** We show that our methods outperform previous approaches with both resource-rich and resource-poor languages. They substantially reduce the gap between random and lexically-overlapping initialization, enabling better model adaption to unseen scripts.

The code for this work is released at [github.com/Adapter-Hub/UNKs\\_everywhere](https://github.com/Adapter-Hub/UNKs_everywhere).

## 2 Background: Multilingual Model Adaptation for Cross-lingual Transfer

Recent language models (Devlin et al., 2019; Conneau et al., 2020a), based on Transformer architectures (Vaswani et al., 2017) and pretrained on massive amounts of multilingual data, have recently

<sup>1</sup>Even languages with unseen scripts share some tokens, e.g. numbers, foreign named entities written in their original scripts, etc. with seen languages.

surpassed (static) cross-lingual word embedding spaces (Ruder et al., 2019; Glavas et al., 2019) as the state-of-the-art paradigm for cross-lingual transfer in NLP (Pires et al., 2019; Wu and Dredze, 2019; Wu et al., 2020; Hu et al., 2020; K et al., 2020). However, recent studies have also indicated that even current state-of-the-art models such as XLM-R (Large) still do not yield reasonable transfer performance across a large number of target languages (Hu et al., 2020). The largest drops are reported for resource-poor target languages (Lauscher et al., 2020), and (even more dramatically) for languages not covered at all during pre-training (Pfeiffer et al., 2020b).

**Standard Cross-Lingual Transfer Setup** with a state-of-the-art pretrained multilingual model such as mBERT or XLM-R is **1)** fine-tuning it on labelled data of a downstream task in a source language and then **2)** applying it directly to perform inference in a target language (Hu et al., 2020). However, as the model must balance between many languages in its representation space, it is not suited to excel at a specific language at inference time without further adaptation (Pfeiffer et al., 2020b).

**Adapters for Cross-lingual Transfer.** Adapter-based approaches have been proposed as a remedy (Rebuffi et al., 2017, 2018; Houlisby et al., 2019; Stickland and Murray, 2019; Bapna and Firat, 2019; Pfeiffer et al., 2020a, 2021). In the cross-lingual setups, the idea is to increase the multilingual model capacity by storing language-specific knowledge of each language in dedicated parameters (Pfeiffer et al., 2020b; Vidoni et al., 2020). We start from MAD-X (Pfeiffer et al., 2020b), a state-of-the-art adapter-based framework for cross-lingual transfer. For completeness, we provide a brief overview of the framework in what follows.

MAD-X comprises three adapter types: language, task, and invertible adapters; this enables learning language and task-specific transformations in a modular and parameter-efficient way. As in prior work (Rebuffi et al., 2017; Houlisby et al., 2019), adapters are trained while keeping the parameters of the pretrained multilingual model fixed. *Language adapters* are trained via masked language modeling (MLM) on unlabelled target language data. *Task adapters* are trained via task-specific objectives on labelled task data in a source language while also keeping the language adapters fixed. Task and language adapters are stacked: this enables the adaptation of the pretrained multilin-

gual model to languages not covered in its pretraining data. MAD-X keeps the same task adapter while substituting the source language adapter with the target language adapter at inference.

In brief, the adapters  $A_l$  at layer  $l$  consist of a down-projection  $\mathbf{D} \in \mathbb{R}^{h \times d}$  where  $h$  is the hidden size of the Transformer model and  $d$  is the dimension of the adapter, followed by a GeLU activation (Hendrycks and Gimpel, 2016) and an up-projection  $\mathbf{U} \in \mathbb{R}^{d \times h}$  at every layer  $l$ :

$$A_l(\mathbf{h}_l, \mathbf{r}_l) = \mathbf{U}_l(\text{GeLU}(\mathbf{D}_l(\mathbf{h}_l))) + \mathbf{r}_l. \quad (1)$$

$\mathbf{h}_l$  and  $\mathbf{r}_l$  are the Transformer hidden state and the residual at layer  $l$ , respectively. The residual connection  $\mathbf{r}_l$  is the output of the Transformer’s feed-forward layer whereas  $\mathbf{h}_l$  is the output of the subsequent layer normalization. For further technical details, we refer the reader to Pfeiffer et al. (2020b).

Current model adaptation approaches (Chau et al., 2020; Wang et al., 2020) generally fine-tune *all* model parameters on target language data. Instead, we follow the more computationally efficient adapter-based paradigm where we keep model parameters fixed, and only train language adapters and target language embeddings. Crucially, while the current adapter-based methods offer extra capacity, they do not offer mechanisms to deal with extended vocabularies of many resource-poor target languages, and do not adapt their representation space towards the target language adequately. This problem is exacerbated when dealing with unseen languages and scripts.<sup>2</sup>

### 3 Cross-lingual Transfer of Lexical Information

The embedding matrix of large multilingual models makes up around 50% of their entire parameter budget (Chung et al., 2021). However, it is not clear how to leverage this large amount of information most effectively for languages that are not adequately represented in the shared multilingual vocabulary due to lack of pretraining data.

A key challenge in using the lexical information encoded in the embedding matrix is to overcome a mismatch in vocabulary between the pretrained model and the target language. To outline this issue, in Table 1 we show for the languages in our NER evaluation the proportion of tokens in each

<sup>2</sup>An alternative approach based on transliteration (Müller et al., 2021) side-steps script adaptation but relies on language-specific heuristics, which are not available for most languages.

Language	iso	Family	Script	% UNKS	% Lex Overl.
English*	en	Indo-Europ.	Latin	0%	66%
Chinese*	zh	Sino-Tibetan	Chinese	0%	79%
Japanese*	ja	Japonic	Japanese	0%	99%
Arabic*	ar	Afro-Asiatic	Arabic	1%	6%
Georgian	ka	Kartvelian	Georgian	2%	27%
Urdu	ur	Indo-Europ.	Arabic	5%	34%
Hindi	hi	Indo-Europ.	Devanagari	2%	33%
Min Dong	cdo	Sino-Tibetan	Chinese	6%	53%
Māori	mi	Austronesian	Latin	1%	45%
Ilokano	ilo	Austronesian	Latin	2%	48%
Guarani	gn	Tupian	Latin	3%	42%
Mingrelian	xmf	Kartvelian	Georgian	7%	22%
Sindhi	sd	Indo-Europ.	Arabic	30%	25%
Erzya	myv	Uralic	Cyrillic	1%	33%
Bhojpuri	bh	Indo-Europ.	Devanagari	1%	28%
Wolof	wo	Niger-Congo	Latin	1%	31%
Amharic	am	Afro-Asiatic	Ge'ez	86%	13%
Tibetan	bo	Sino-Tibetan	Tibetan	66%	20%
Khmer	km	Austroasiatic	Khmer	79%	19%
Divehi	dv	Indo-Europ.	Thaana	85%	14%
Sinhala	si	Indo-Europ.	Sinhala	75%	23%

Table 1: Languages in our NER evaluation, together with their language family and common script. For each monolingual vocabulary (§4.3), we compute the proportion of tokens that cannot be composed by the subword tokens from the original mBERT vocabulary (UNKs) as well as of *lexically overlapping* tokens that are present in both vocabularies. \*: For the high-resource source languages, proportions are calculated with regard to the tokenizers used by Rust et al. (2021).

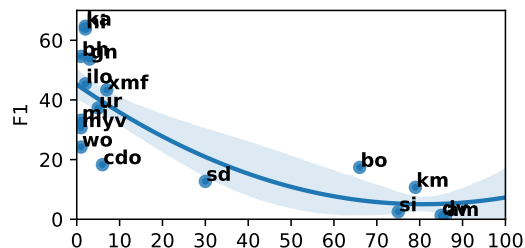


Figure 2: mBERT’s zero-shot cross-lingual transfer performance with respect to the proportion of UNKS in mBERT’s original vocabulary relative to the target language (see Table 1).

language that are effectively *unknown* (UNK) to mBERT: they occur in the vocabulary of a separately trained monolingual tokenizer (§4.3), but cannot even be composed by subword tokens from the original mBERT’s vocabulary. Table 1 also provides the proportion of *lexically overlapping* tokens, i.e., tokens that are present both in mBERT’s and monolingual in-language vocabularies. The zero-shot performance of mBERT generally deteriorates with less lexical overlap and more UNKS in a target language: see Figure 2. Pearson’s  $\rho$  correlation scores between the lexical overlap and proportion of UNKS (see Table 1) and NER performance are 0.443 and  $-0.798$ , respectively.

Method	Special tokens	Lexical overlap	Latent semantic concepts	Language clusters	New params	# of new params	Reference
EL-RAND	✓				$\mathbf{X}'$	7.68M	Artetxe et al. (2020)
EL-LEX	✓	✓			$\mathbf{X}'$	7.68M	Chau et al. (2020); Wang et al. (2020)
$\text{MF}_*^C\text{-RAND}$	✓		✓	✓	$\mathbf{F}', \mathbf{I}'$	$1\text{M} + C \cdot 10\text{k}$	Ours
$\text{MF}_*^C\text{-LEX}$	✓	✓	✓	✓	$\mathbf{F}', \mathbf{I}'$	$1\text{M} + C \cdot 10\text{k}$	Ours

Table 2: Overview of our methods and related approaches together with the pretrained knowledge they utilize. We calculate the number of new parameters per language with  $V' = 10k$ ,  $D = 768$ , and  $D' = 100$ . We do not include up-projection matrices  $\mathbf{G}$  as these are learned only once and make up a comparatively small number of parameters.

Recent approaches such as invertible adapters (Pfeiffer et al., 2020b) that adapt embeddings in the pretrained multilingual vocabulary may be able to deal with lesser degrees of lexical overlap. Still, they cannot deal with UNK tokens. In the following, we systematize existing approaches and present novel ways of adapting a pretrained model to the vocabulary of a target language that can handle this challenging setting, present most acutely when adapting to languages with unseen scripts. We summarize the approaches in Table 2 based on what types of pretrained information they utilize. All approaches rely on a new vocabulary  $V'$ , learned on the target language data.

### 3.1 Target-Language Embedding Learning

A straightforward way to adapt a pretrained model to a new language is to learn new embeddings for the language. Given the new vocabulary  $V'$ , we initialize new embeddings  $\mathbf{X}' \in \mathbb{R}^{|V'| \times D}$  for all  $V'$  vocabulary items where  $D$  is the dimensionality of the existing embeddings  $\mathbf{X} \in \mathbb{R}^{|V| \times D}$ , and only initialize special tokens (e.g. [CLS], [SEP]) with their pretrained representations. We train the new embeddings of the  $\mathbf{X}'$  with the pretraining task. This approach, termed EL-RAND, was proposed by Artetxe et al. (2020): they show that it allows learning aligned representations for a new language but only evaluate on high-resource languages. The shared special tokens allow the model to access a minimum amount of lexical information, which can be useful for transfer (Dufter and Schütze, 2020). Beyond this, this approach leverages knowledge from the existing embedding matrix only implicitly to the extent that the higher-level hidden representations are aligned to the lexical representations.

### 3.2 Initialization with Lexical Overlap

To leverage more lexical information, we can apply shared initialization not only to the special tokens but to all lexically overlapping tokens. Let

us denote this vocabulary subset with  $V'_{lex}$ , and  $V'_{rand} = V' \setminus V'_{lex}$ . In particular, we initialize the embeddings of all lexically overlapping tokens  $\mathbf{X}'_{lex}$  from  $V'_{lex}$  with their pretrained representations from the original matrix  $\mathbf{X}$ , while the tokens from  $V'_{rand}$  receive randomly initialized embeddings  $\mathbf{X}'_{rand}$ . We then fine-tune all target language embeddings  $\mathbf{X}' = \mathbf{X}'_{lex} \cup \mathbf{X}'_{rand}$  on the target language data. Wang et al. (2020) cast this as extending  $V$  with new tokens. In contrast, we seek to disentangle the impact of vocabulary size and pretrained information. As one variant of this approach, Chau et al. (2020) only add the 99 most common tokens of a new language to  $V$ .

Initialization with lexical overlap, termed EL-LEX, allows us to selectively leverage the information from the pretrained model on a per-token level based on surface-level similarity. Intuitively, this should be most useful for languages that are lexically similar to those seen during pretraining and have a substantial proportion of lexically overlapping tokens. However, such lexical overlap is a lot rarer for languages that are written in different scripts. For such languages, relying on surface-level string similarity alone may not be enough.

### 3.3 Embedding Matrix Factorization

We therefore propose to identify latent semantic concepts in the pretrained model’s embedding matrix that are general across languages and useful for transfer. Further, to allow modeling flexibility we propose to learn a grouping of similar tokens. We achieve this by factorizing the pretrained embedding matrix  $\mathbf{X} \in \mathbb{R}^{|V| \times D}$  into lower-dimensional word embeddings  $\mathbf{F} \in \mathbb{R}^{|V| \times D'}$  and  $C$  shared up-projections  $\mathbf{G}^1, \dots, \mathbf{G}^C \in \mathbb{R}^{D' \times D}$  that encode general cross-lingual information:

$$\mathbf{X} \approx \sum_{c \in C} \text{diag}(\mathbf{i}_c) \mathbf{F} \mathbf{G}^c \quad (2)$$

$D'$  is the dimensionality of the lower-dimensional embeddings.  $\mathbf{I} = [\mathbf{i}_1, \dots, \mathbf{i}_C] \in \mathbb{R}^{|V| \times C}$  is an

indicator matrix where  $i_{v,c} = 1$  iff token  $v$  is associated with the  $c$ -th up-projection, else 0.<sup>3</sup>  $\text{diag}(\cdot)$  creates a diagonal matrix from a vector, i.e.  $\text{diag}(\mathbf{i}_c) \in \mathbb{R}^{|V| \times |V|}$ .

**MF<sup>1-\*</sup>**. In the basic case where  $C = 1$ , Eq. (2) simplifies to  $\mathbf{X} \approx \mathbf{F}\mathbf{G}$ . As  $\mathbf{X}$  is unconstrained, we follow Ding et al. (2008) and interpret this as a semi-non-negative matrix factorization (Semi-NMF) problem. In Semi-NMF,  $\mathbf{G}$  is restricted to be non-negative while no restrictions are placed on the signs of  $\mathbf{F}$ . The Semi-NMF is computed via an iterative updating algorithm that alternatively updates  $\mathbf{F}$  and  $\mathbf{G}$  where the Frobenius norm is minimized.<sup>4</sup>

$\mathbf{G}$  is shared across all tokens and thus encodes general properties of the original embedding matrix  $\mathbf{X}$  whereas  $\mathbf{F}$  stores token-specific information.  $\mathbf{G}$  only needs to be pretrained once and can be used and fine-tuned for every new language. To this end, we simply learn new low-dimensional embeddings  $\mathbf{F}' \in \mathbb{R}^{|V'| \times D'}$  with the pretraining task, which are up-projected with  $\mathbf{G}$  and fed to the model.

**MF<sub>KMEANS</sub><sup>C-\*</sup>**. When  $C > 1$ , each token is associated with one of  $C$  up-projection matrices. Grouping tokens and using a separate up-projection matrix per group may help balance sharing information across typologically similar languages with learning a robust representation for each token (Chung et al., 2020). We propose two approaches to automatically learn such a clustering.

In our first, pipeline-based approach, we first cluster  $\mathbf{X}$  into  $C$  clusters using KMeans. For each cluster, we then factorize the subset of embeddings  $\mathbf{X}^c$  associated with the  $c$ -th cluster separately using Semi-NMF equivalently as for MF<sup>1-\*</sup>.

For a new language, we learn new low-dim embeddings  $\mathbf{F}' \in \mathbb{R}^{|V'| \times D'}$  and a randomly initialized matrix  $\mathbf{Z} \in \mathbb{R}^{|V'| \times C}$ , which allows us to compute the cluster assignment matrix  $\mathbf{I}' \in \mathbb{R}^{|V'| \times C}$ . Specifically, for token  $v$ , we obtain its cluster assignment as  $\arg \max$  of  $\mathbf{z}'_{v,\cdot}$ . As  $\arg \max$  is not differentiable, we employ the Straight-Through Gumbel-Softmax estimator (Jang et al., 2017) defined as:

$$i'_{v,c} = \frac{\exp(\log(z_{v,c}) + g_c)/\tau)}{\sum_{j=1}^C \exp(\log(z_{v,j}) + g_j)/\tau}, \quad (3)$$

where  $\tau$  is a temperature parameter, and  $\mathbf{g} \in \mathbb{R}^{|V|}$  corresponds to samples from the Gumbel distri-

<sup>3</sup>We use bold upper case letters ( $\mathbf{X}$ ) for matrices, subscripts with bold letters ( $\mathbf{x}_i$ ) for rows or columns and subscripts with standard weight letters for specific elements ( $x_i, x_{i,j}$ ).

<sup>4</sup>For more details see Ding et al. (2008).

bution  $g_j \sim -\log(-\log(u_j))$  with  $u_j \sim \mathcal{U}(0, 1)$  being the uniform distribution.  $\mathbf{z}_{v,\cdot}$  can be seen as “logits” used for assigning the  $v$ -th token a cluster. As  $\tau \rightarrow 0$ , the softmax becomes an  $\arg \max$  and the Gumbel-Softmax distribution approximates more closely the categorical distribution.  $\mathbf{I}' \in \mathbb{R}^{|V'| \times C}$  represents the one-hot encoded, indicator function over possible clusters, with learnable parameters  $\mathbf{Z}$ . As before,  $i'_{v,c} = 1$  iff new token  $v$  is associated with up-projection  $c$ , else 0.

**MF<sub>NEURAL</sub><sup>C-\*</sup>**. We can also learn the cluster assignment and up-projections jointly. Specifically, we parameterize  $\mathbf{G}$  in Eq. (2) using a neural net where we learn the indicator matrix  $\mathbf{I}$  equivalently to Eq. (3). The objective minimizes the L<sub>2</sub>-norm between the original and predicted embeddings:

$$\mathcal{L} = \|\mathbf{X} - \sum_{c \in C} \text{diag}(\mathbf{i}_c)\mathbf{F}\mathbf{G}^c\|_2 \quad (4)$$

For a new language, we proceed analogously.

**MF<sup>\*</sup>-RAND and MF<sup>\*</sup>-LEX**. Finally, we can combine different initialization strategies (see §3.1 and §3.2) with the embedding matrix factorization technique. We label the variant which relies on random initialization, see §3.1, as MF<sup>1</sup>-RAND. The variant, which relies on lexically overlapping tokens from §3.2 can leverage both surface-level similarity as well as latent knowledge in the embedding matrix; we simply initialize the embeddings of overlapping lexical tokens (from  $V'_{lex}$ ) in  $\mathbf{F}'$  with their low-dim representations from  $\mathbf{F}$ . The remaining tokens (from  $V'_{rand}$ ) are randomly initialized in  $\mathbf{F}'$ .

Factorizing the embedding matrix has the additional benefit of reducing the number of trainable parameters and correspondingly the amount of storage space required for each additional language. This is especially true when  $D \gg D'$ .<sup>5</sup>

## 4 Experiments

**Data.** For pretraining, we leverage the Wikipedia dumps of the target languages. We conduct experiments on named entity recognition (NER) and dependency parsing (DP). For NER, we use the WikiAnn (Pan et al., 2017) dataset, partitioned into train, dev, and test portions by Rahimi et al. (2019). For DP we use Universal Dependencies (UD; Nivre et al., 2016, 2020; Zeman et al., 2020).

<sup>5</sup>For instance, when  $|V| = |V'| = 10k$ ,  $D = 768$ , and  $D' = 100$  as in our experiments we reduce the amount of storage space required by 82% per language.

**Languages.** WikiAnn offers a wide language coverage (176 languages) and, consequently, a number of language-related comparative analyses. In order to systematically compare against state-of-the-art cross-lingual methods under different evaluation conditions, we identify **1)** low-resource languages where the script has been covered by mBERT but the model has not been specifically trained on the language **2)** as well as low-resource languages with scripts not covered at all by pretraining. In each case, we select four languages, taking into account variance in data availability and typological diversity. We select four high-resource source languages (English, Chinese, Japanese, Arabic) in order to go beyond English-centric transfer. We evaluate the cross-lingual transfer performance of these 4 languages to the 17 diverse languages. For DP, we chose the subset that occurs in UD. We highlight the properties of all 21 languages in Table 1.

#### 4.1 Baselines

**mBERT (Standard Transfer Setup).** We primarily focus on mBERT as it has been shown to work well for low-resource languages (Pfeiffer et al., 2020b). mBERT is trained on the 104 languages with the largest Wikipedias.<sup>6</sup> In the standard cross-lingual transfer setting (see §2), the full model is fine-tuned on the target task in the (high-resource) source language, and is evaluated on the test set of the target (low-resource) language.

**MAD-X.** We follow Pfeiffer et al. (2020b) and stack task adapters on top of pretrained language adapters (see §2). When training the model on source language task data, only the task adapter is trained while the original model weights and the source language adapter are frozen. At inference, the source language adapter is replaced with the target language adapter.

**MAD-X 2.0.** The adapter in the last transformer layer is not encapsulated between frozen transformer layers, and can thus be considered an extension of the prediction head. This places no constraints on the representation of the final adapter, possibly decreasing transfer performance when replacing the language adapters for zero-shot transfer. In this work, we thus propose to drop the adapter in the last transformer layer, and also evaluate this novel variant of the MAD-X framework.

<sup>6</sup>See Appendix Table 12 for the list of all 104 covered languages with corresponding scripts.

#### 4.2 Methods

We experiment with the methods from Table 2 and discussed in §3, summarized here for clarity.

**EL-\*** We randomly initialize embeddings for all tokens—except special tokens—in the new vocabulary (EL-RAND) or initialize embeddings of lexically overlapping tokens with their pretrained representations (EL-LEX).

**MF<sup>1</sup>-\*** We randomly initialize lower-dimensional embeddings (MF<sup>1</sup>-RAND) or initialize lexically overlapping tokens with their corresponding lower-dimensional pretrained representation (MF<sup>1</sup>-LEX) while using a single pretrained projection matrix.

**MF<sub>\*</sub><sup>10</sup>-\*** We learn assignments to 10 clusters via k-means and up-projection matrices via Semi-NMF (MF<sub>KMEANS</sub><sup>10</sup>-\*). Alternatively, we learn cluster assignments with Gumbel-Softmax and up-projection matrices jointly (MF<sub>NEURAL</sub><sup>10</sup>-\*). For new tokens we use random (MF<sub>\*</sub><sup>10</sup>-RAND) or lexical overlap initialisation (MF<sub>\*</sub><sup>10</sup>-LEX).

#### 4.3 Experimental Setup

Previous work generally fine-tunes the entire model on the target task (Chau et al., 2020; Wang et al., 2020). To extend the model to a new vocabulary, Artetxe et al. (2020) alternately freeze and fine-tune embeddings and transformer weights for pretraining, and target task fine-tuning, respectively. We find that this approach largely underperforms adapter-based transfer as proposed by Pfeiffer et al. (2020b), and we thus primarily focus on adapter-based training in this work.<sup>7</sup>

**Adapter-Based Transfer.** We largely follow the experimental setup of Pfeiffer et al. (2020b), unless noted otherwise. We obtain language adapters for the high-resource languages from [AdapterHub.ml](https://adapterhub.ml) (Pfeiffer et al., 2020a) and train language adapters and embeddings for the low-resource languages jointly while keeping the rest of the model fixed. For zero-shot transfer, we replace the source language adapter with the target adapter, and also replace the entire embedding layer with the new embedding layer specialized to the target language. MAD-X 2.0 consistently outperforms MAD-X (see §5); we thus use this setup for all our methods.

**Tokenizer.** We learn a new WordPiece tokenizer for each target language with a vocabulary size of 10k using the HuggingFace tokenizer library.<sup>8</sup>

<sup>7</sup>We present results with such full model transfer in §A.1.

<sup>8</sup><https://github.com/huggingface/tokenizers>

	Seen Languages				Unseen Languages but Covered Scripts										New Scripts				<i>Mac. Avg</i>		
	ka	ur	hi	Avg	cdo	mi	ilo	gn	xmf	sd	myv	bh	wo	Avg	am	bo	km	dv		si	Avg
mBERT	64.7	37.3	<b>63.8</b>	55.3	18.3	33.2	45.4	53.7	43.4	12.7	30.7	54.6	24.2	35.1	0.9	17.4	10.7	1.3	2.5	6.6	<i>30.3</i>
MAD-X	<b>63.9</b>	51.4	58.7	<b>58.0</b>	29.1	45.6	45.5	52.7	51.9	34.0	57.8	57.0	49.5	47.0	<b>10.8</b>	<b>24.8</b>	17.6	16.8	16.8	17.4	<i>40.2</i>
MAD-X 2.0	<b>65.6</b>	<b>55.3</b>	61.0	<b>60.6</b>	<b>30.7</b>	<b>50.6</b>	<b>64.0</b>	<b>56.3</b>	<b>52.6</b>	<b>37.1</b>	<b>63.0</b>	<b>59.8</b>	<b>55.5</b>	<b>52.2</b>	10.7	24.7	<b>18.1</b>	<b>22.2</b>	<b>18.7</b>	<b>17.9</b>	<i>43.9</i>
EL-RAND	<b>65.8</b>	47.8	<b>63.8</b>	<b>59.1</b>	38.0	7.3	<b>57.9</b>	48.5	59.4	44.2	35.2	55.5	5.1	39.0	<b>42.9</b>	<b>53.9</b>	59.5	32.7	51.5	48.1	<i>45.2</i>
MF <sup>1</sup> -RAND	63.5	48.5	57.8	56.6	39.8	<b>19.3</b>	43.6	47.3	57.9	45.1	<b>60.4</b>	44.4	<b>45.0</b>	<b>44.8</b>	42.6	49.0	61.7	<b>43.4</b>	<b>52.1</b>	<b>49.8</b>	<i>48.3</i>
MF <sup>10</sup> <sub>KMEANS</sub> -RAND	64.3	<b>52.6</b>	60.4	<b>59.1</b>	20.7	11.7	51.2	50.8	58.2	<b>45.2</b>	49.0	<b>60.0</b>	42.5	43.2	37.3	37.5	<b>64.2</b>	18.5	47.2	40.9	<i>45.4</i>
MF <sup>10</sup> <sub>NEURAL</sub> -RAND	63.2	49.4	60.6	57.7	<b>40.7</b>	5.5	48.5	<b>52.0</b>	<b>60.4</b>	27.2	46.7	54.5	41.2	41.9	42.7	38.4	63.7	38.7	48.0	46.3	<i>46.0</i>
EL-LEX	<b>69.3</b>	57.9	<b>66.4</b>	<b>64.5</b>	46.8	32.2	<b>58.8</b>	54.1	57.0	44.6	<b>56.3</b>	59.5	50.1	51.0	<b>46.5</b>	51.5	61.0	<b>47.2</b>	<b>56.3</b>	<b>52.5</b>	<i>53.6</i>
MF <sup>1</sup> -LEX	65.5	54.0	61.5	60.4	46.6	39.8	55.2	<b>54.8</b>	<b>57.2</b>	45.0	55.0	59.1	48.1	51.2	38.6	42.5	63.9	43.8	51.5	48.0	<i>51.2</i>
MF <sup>10</sup> <sub>KMEANS</sub> -LEX	66.3	<b>59.5</b>	61.1	62.3	<b>48.9</b>	47.5	46.4	53.6	56.7	46.2	58.4	<b>61.0</b>	<b>57.9</b>	<b>53.0</b>	45.5	<b>51.6</b>	64.2	44.9	46.2	50.5	<i>53.9</i>
MF <sup>10</sup> <sub>NEURAL</sub> -LEX	67.0	55.8	62.6	61.8	47.8	<b>50.8</b>	53.1	53.9	55.6	<b>46.4</b>	50.4	60.6	51.8	52.3	46.7	43.0	<b>65.2</b>	46.3	51.5	50.5	<i>53.4</i>

(a) Named Entity Recognition: Mean  $F_1$  test results for UD averaged over 5 runs and averaged over the 4 source languages.

	Seen Languages			Unseen Languages but Covered Scripts				New Script		<i>Macro Avg</i>
	hi	ur	Avg	bh	myv	wo	Avg	am	UAS / LAS	
mBERT	<b>45.8 / 29.5</b>	34.9 / 20.5	<b>40.3 / 25.0</b>	<b>33.9 / 18.6</b>	29.7 / 13.0	27.5 / 7.5	30.4 / 13.0	10.1 / 3.8	<i>30.3 / 15.5</i>	
MAD-X 2.0	42.0 / 26.1	<b>35.0 / 20.1</b>	38.5 / 23.1	31.0 / 16.2	<b>47.6 / 29.8</b>	<b>35.4 / 18.4</b>	<b>38.0 / 21.4</b>	<b>14.0 / 7.7</b>	<i>34.2 / 19.7</i>	
EL-RAND	41.9 / 25.9	32.9 / 18.4	37.4 / 22.1	29.9 / 13.8	45.4 / 23.9	24.3 / 7.2	33.2 / 15.0	30.1 / 11.3	<i>34.1 / 16.7</i>	
MF <sup>1</sup> -RAND	41.7 / 26.2	33.7 / 19.4	37.7 / 22.8	28.7 / 13.7	47.6 / 27.2	<b>32.7 / 15.5</b>	36.3 / 18.8	<b>34.3 / 13.3</b>	<i>36.4 / 19.2</i>	
MF <sup>10</sup> <sub>KMEANS</sub> -RAND	40.7 / 25.1	30.4 / 17.7	35.6 / 21.4	<b>31.9 / 16.3</b>	47.8 / 27.7	29.6 / 13.7	<b>36.4 / 19.2</b>	22.6 / 10.1	<i>33.8 / 18.4</i>	
MF <sup>10</sup> <sub>NEURAL</sub> -RAND	<b>43.1 / 27.0</b>	<b>34.7 / 20.0</b>	<b>38.9 / 23.5</b>	30.8 / 14.7	<b>48.6 / 28.7</b>	21.7 / 9.4	33.7 / 17.6	32.4 / 12.4	<i>35.2 / 18.7</i>	
EL-LEX	41.9 / 26.1	34.0 / 19.6	37.9 / 22.9	30.6 / 16.2	48.8 / 27.7	37.9 / 19.2	39.1 / 21.0	34.0 / 9.9	<i>36.6 / 19.8</i>	
MF <sup>1</sup> -LEX	42.1 / 26.7	34.0 / 20.1	38.0 / 23.4	<b>31.2 / 16.4</b>	<b>49.0 / 28.9</b>	37.6 / 19.5	<b>39.3 / 21.6</b>	<b>34.8 / 12.9</b>	<i>36.8 / 20.8</i>	
MF <sup>10</sup> <sub>KMEANS</sub> -LEX	42.2 / 27.0	34.7 / 20.3	38.5 / 23.6	31.1 / <b>16.4</b>	47.7 / 28.6	33.8 / 17.3	37.5 / 20.8	32.2 / 13.2	<i>37.0 / 20.5</i>	
MF <sup>10</sup> <sub>NEURAL</sub> -LEX	<b>43.0 / 27.4</b>	<b>36.1 / 21.2</b>	<b>39.6 / 24.3</b>	30.9 / 16.2	46.8 / 27.6	<b>38.7 / 20.4</b>	38.8 / 21.4	31.4 / 12.1	<i>37.8 / 20.8</i>	

(b) Dependency Parsing: Mean UAS and LAS test results for UD averaged over 5 runs and averaged over the 4 source languages.

Table 3: Mean test results for (a) NER and (b) DP. The top group (first two/three rows) includes models, which leverage the original tokenizer which is not specialized for the target language. The second group (last eight rows) include models with new tokenizers. We separate models with randomly initialized embeddings (\*-RANDINIT) from models with lexical init (\*-LEXINIT) by the dashed line. Bold numbers indicate best-scoring models of the respective group, underlined numbers the best performance overall. Source languages are *en*, *ar*, *zh*, and *ja*.

**Semi-NMF.** We factorize the pretrained embedding matrix of mBERT using Semi-NMF (Ding et al., 2008) leveraging the default implementation provided by Bauckhage et al. (2011).<sup>9</sup> We train for 3,000 update steps and leverage the corresponding matrices  $\mathbf{F}$  and  $\mathbf{G}$  as initialization for the new vocabulary. We choose the reduced embedding dimensionality  $D' = 100$ .  $\mathbf{F}$  is only used when initializing the (lower-dimensional) embedding matrix with lexically overlapping representations.

**Masked Language Modeling.** For MLM pretraining we leverage the entire Wikipedia corpus of the respective language. We train for 200 epochs or  $\sim 100k$  update steps, depending on the corpus size. The batch size is 64; the learning rate is  $1e - 4$ .

**Task Fine-tuning.** Our preliminary experiments suggested that fine-tuning the model for a smaller number of epochs leads to better transfer performance in low-resource languages in general. We thus fine-tune all the models for 10 epochs, evaluating on the source language dev set after every epoch. We then take the best-performing model according to the dev  $F_1$  score, and use it in zero-shot

transfer. We train all the models with a batch size of 16 on high resource languages. For NER we use learning rates  $2e - 5$  and  $1e - 4$  for full fine-tuning and adapter-based training, respectively. For DP, we use a transformer-based variant (Glavas and Vulic, 2021) of the standard deep biaffine attention dependency parser (Dozat and Manning, 2017) and train with learning rates  $2e - 5$  and  $5e - 4$  for full fine-tuning and adapter-based training respectively.

## 5 Results and Discussion

The main results are summarised in Table 3a for NER, and in Table 3b for DP. First, our novel MAD-X 2.0 considerably outperforms the MAD-X version of Pfeiffer et al. (2020b). However, while both MAD-X versions improve over mBERT for unseen scripts, the performance remains quite low on average. The corresponding mBERT tokenizer is not able to adequately represent unseen scripts: many tokens are substituted by UNKS (see Table 1), culminating in the observed low performance.

For our approaches that learn new embedding matrices, we observe that for languages seen during pretraining, but potentially underrepresented by the

<sup>9</sup><https://github.com/cthurau/pymf>

	ka	ur	hi	cdo	mi	ilo	gn	xmf	sd	myv	bh	wo	am	bo	km	dv	si
Numbers	11%	8%	6%	26%	17%	8%	8%	15%	9%	11%	11%	7%	13%	9%	3%	12%	9%
Lat Char	4%	3%	3%	2%	2%	2%	2%	5%	4%	3%	4%	3%	8%	5%	5%	7%	4%
Lat (S)W	10%	9%	7%	54%	61%	64%	65%	10%	10%	9%	26%	71%	19%	30%	31%	17%	30%
Oth Char	55%	36%	48%	18%	19%	26%	25%	60%	39%	26%	24%	18%	60%	52%	61%	63%	57%
Oth (S)W	20%	43%	36%	0%	1%	0%	0%	10%	38%	51%	35%	1%	0%	3%	0%	1%	0%

Table 4: Grouping of tokens that lexically overlap between the original mBERT tokenizer and the tokenizer of the target language. *Numbers* includes all tokens which include at least one number; *Lat Char* indicates all Latin tokens of length 1; *Lat (S)W* includes all (sub)words that include a Latin character but are of length  $> 1$ . Consequently, *Oth Char* and *Oth (S)W* consists of characters and (sub)words respectively, which do not include Latin characters.

cdo	mi	ilo	gn	bo	km	dv	si
Northumberland	Massachusetts	Melastomataceae	establecimiento	University	languages	government	International
Massachusetts	Encyclopedia	munisipalidad	vicepresidente	Therefore	language	Chinese	Bangladesh
International	Pennsylvania	Internasional	Internacional	suffering	formula	govern	wikipedia
Pennsylvania	Jacksonville	internasional	internacional	existence	disease	system	Australia
Philadelphia	Turkmenistan	International	Independencia	practice	control	ation	Zimbabwe

Table 5: Longest lexically overlapping (sub)words.

model (e.g. Georgian (ka), Urdu (ur), and Hindi (hi)), the proposed methods outperform MAD-X 2.0 for all tasks. This is in line with contemporary work (Rust et al., 2021), which emphasizes the importance of tokenizer quality for the downstream task. Consequently, for unseen languages with under-represented scripts, the performance gains are even larger, e.g., we see large improvements for Min Dong (cdo), Mingrelian (xmf), and Sindhi (sd). For unseen languages with the Latin script, our methods perform competitively (e.g. Maori (mi), Ilokano (ilo), Guarani (gn), and Wolof (wo)): this empirically confirms that the Latin script is adequately represented in the original vocabulary. The largest gains are achieved for languages with unseen scripts (e.g. Amharic (am), Tibetan (bo), Khmer (km), Divehi (dv), Sinhala (si)), as these languages are primarily represented as UNK tokens by the original mBERT tokenizer.

We observe improvements for most languages with lexical overlap initialization. This adds further context to prior studies which found that a shared vocabulary is not necessary for learning multilingual representations (Conneau et al., 2020b; Artetxe et al., 2020): while it is possible to generalize to new languages without lexical overlap, leveraging the overlap still offers additional gains.

The methods based on matrix factorization ( $MF_{*-*}^*$ ) improve performance over full-sized embedding methods ( $EL_{*-*}^*$ ), especially in the setting without lexical overlap initialization ( $*-RAND$ ). This indicates that by factorizing the information encoded in the original embedding matrix we are

able to extract relevant information for unseen languages. Combining matrix factorization with lexical overlap initialization ( $MF_{*-*}^*-LEX$ ), zero-shot performance improves further for unseen languages with covered scripts. This suggests that the two methods complement each other. For 6/9 of these languages, we find that encoding the embeddings in multiple up-projections ( $MF_{*-*}^{10}$ ) achieves the peak score. This in turn verifies that grouping similar tokens improves the robustness of token representations (Chung et al., 2020). For unseen languages with covered scripts, this model variant also outperforms MAD-X 2.0 on average.

For languages with unseen scripts we find that MF has smaller impact. While the encoded information supports languages similar to those seen by the model in pretraining, languages with unseen scripts are too distant to benefit from this latent multilingual knowledge. Surprisingly, lexical overlap is helpful for languages with unseen scripts.

Overall, we observe that both  $MF_{KMEANS}^{10}$  and  $MF_{NEURAL}^{10}$  perform well for most languages, where the KMEANS variant performs better for NER and the NEURAL variant performs better for UD.

## 6 Further Analysis

### 6.1 Lexically Overlapping (Sub)Words

We perform a quantitative analysis of lexically overlapping tokens, i.e. that occur both in mBERT’s and monolingual in-language vocabularies; see Table 4. For languages with scripts not covered by the mBERT tokenizer, most lexically-overlapping tokens are single characters of a non-Latin script.



	Seen Languages				Unseen Languages but Covered Scripts								New Scripts				
	ka	ur	hi	cdo	mi	ilo	gn	xmf	sd	myv	bh	wo	am	bo	km	dv	si
MF <sub>KMEANS</sub> <sup>10</sup> -LEX 100	66.3	59.5	61.1	48.9	47.5	46.4	53.6	56.7	46.2	58.4	61.0	57.9	45.5	51.6	64.2	44.9	46.2
MF <sub>KMEANS</sub> <sup>10</sup> -LEX 300	65.1	46.8	56.8	26.6	5.5	53.7	10.0	52.2	43.5	37.4	53.2	3.7	21.5	7.4	63.9	38.8	51.2

Table 6: Mean  $F_1$  test results averaged over 5 runs and the 4 high-resource source languages English, Chinese, Japanese, and Arabic. The first row presents results with 100-, the second row 300-dimensional embeddings.

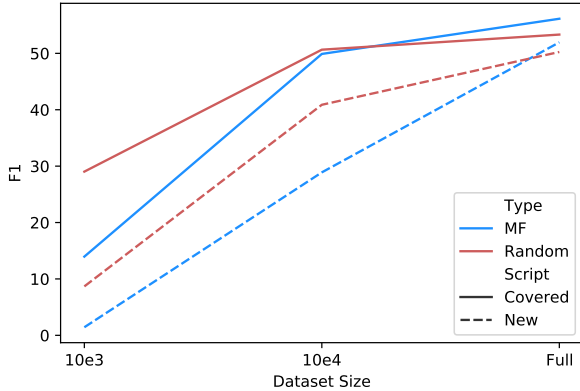


Figure 3: Sample efficiency. For "MF" we leverage the MF<sub>KMEANS</sub><sup>10</sup>-LEX setting.

We further present the top 5 longest lexically overlapping (sub) words for four languages with scripts covered during pretraining (Min Dong, Maori, Ilokano, and Guarani) and four languages with unseen scripts (Tibetan, Khmer, Divehi, Sinhala) in Table 5. We observe that frequent lexically overlapping tokens are named entities in the Latin script, indicating that NER may not be the best evaluation task to objectively assess generalization performance to such languages. If the same named entities also occur in the training data of higher-resource languages, the models will be more successful at identifying them in the unseen language, which belies a lack of deeper understanding of the low-resource language. This might also explain why greater performance gains were achieved for NER than for DP.

## 6.2 Sample Efficiency

We have further analyzed the sample efficiency of our approaches. We find that EL-LEX is slightly more sample-efficient than MF<sub>KMEANS</sub><sup>10</sup>-LEX, where the latter outperforms the former with more data available. In Figure 3 we plot the zero-shot transfer performance where the adapters and embeddings were pretrained on different amounts of data. We further find that lower-dimensional embeddings tend to outperform higher-dimensional embeddings for the majority of languages. In Table 6 we compare 300 with 100 dimensional embeddings for the MF<sub>KMEANS</sub><sup>10</sup>-LEX approach.

## 6.3 Script Clusters

We analyzed the KMeans clusters based on the tokens that consist of characters of a certain script in Figure 4 of the Appendix. We find distinct script-based groups; For instance, 5 clusters consist primarily of Latin-script tokens<sup>10</sup>, two clusters predominantly consist of Chinese, and a few Korean tokens. Interestingly, 2 clusters consisted of Cyrillic, and Arabic scripts as well as scripts used predominantly in India, varying slightly in their distribution. Lastly, one cluster included tokens of all except the Latin script.

## 7 Conclusion

We have systematically evaluated strategies for model adaptation to unseen languages with seen and unseen scripts. We have assessed the importance of the information stored within the original embedding matrix by means of leveraging lexically overlapping tokens, and extracting latent semantic concepts. For the latter, we have proposed a new method of encoding the embedding matrix into lower-dimensional embeddings and up-projections. We have demonstrated that our methods outperform previous approaches on NER and dependency parsing for both resource-rich and resource-poor scenarios, reducing the gap between random and lexical overlap initialisation, and enabling more effective model adaptation to unseen scripts.

## Acknowledgments

Jonas Pfeiffer is supported by the LOEWE initiative (Hesse, Germany) within the emergenCITY center. The work of Ivan Vulić is supported by the ERC Grant LEXICAL (no. 648909) and the ERC PoC Grant MultiConvAI (no. 957356).

We thank Laura Rimell, Nils Reimers, Michael Bugert and the anonymous reviewers for insightful feedback and suggestions on a draft of this paper.

<sup>10</sup>The majority of tokens of mBERT use the Latin script.

## References

- Judit Ács. 2019. [Exploring BERT’s Vocabulary](#). *Blog Post*.
- Alan Ansell, Edoardo Maria Ponti, Jonas Pfeiffer, Sebastian Ruder, Goran Glavaš, Ivan Vulić, and Anna Korhonen. 2021. [MAD-G: Multilingual Adapter Generation for Efficient Cross-Lingual Transfer](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, Online.
- Mikel Artetxe, Sebastian Ruder, and Dani Yogatama. 2020. [On the cross-lingual transferability of monolingual representations](#). In *Proceedings of the 58th Conference of the Association for Computational Linguistics, ACL 2020, Virtual Conference, July 6-8, 2020*, pages 4623–4637.
- Ankur Bapna and Orhan Firat. 2019. [Simple, scalable adaptation for neural machine translation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 1538–1548.
- Christian Bauckhage, Kristian Kersting, and Christian Thureau. 2011. [Factorizing gigantic matrices](#). *ECML-PKDD 2011: Tutorial*.
- Ethan C. Chau, Lucy H. Lin, and Noah A. Smith. 2020. [Parsing with multilingual bert, a small treebank, and a small corpus](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings, EMNLP 2020, Online Event, 16-20 November 2020*, pages 1324–1334.
- Hyung Won Chung, Thibault Févry, Henry Tsai, Melvin Johnson, and Sebastian Ruder. 2021. [Re-thinking Embedding Coupling in Pre-trained Language Models](#). In *9th International Conference on Learning Representations, ICLR 2021, Online, 2021, Conference Track Proceedings*.
- Hyung Won Chung, Dan Garrette, Kiat Chuan Tan, and Jason Riesa. 2020. [Improving multilingual models with language-clustered vocabularies](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 4536–4546.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020a. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Conference of the Association for Computational Linguistics, ACL 2020, Virtual Conference, July 6-8, 2020*, pages 8440–8451.
- Alexis Conneau, Shijie Wu, Haoran Li, Luke Zettlemoyer, and Veselin Stoyanov. 2020b. [Emerging cross-lingual structure in pretrained language models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6022–6034.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186.
- Chris HQ Ding, Tao Li, and Michael I Jordan. 2008. [Convex and semi-nonnegative matrix factorizations](#). *IEEE transactions on pattern analysis and machine intelligence*, 32(1):45–55.
- Timothy Dozat and Christopher D. Manning. 2017. [Deep biaffine attention for neural dependency parsing](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Philipp Dufter and Hinrich Schütze. 2020. [Identifying elements essential for bert’s multilinguality](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 4423–4437.
- Goran Glavas, Robert Litschko, Sebastian Ruder, and Ivan Vulić. 2019. [How to \(properly\) evaluate cross-lingual word embeddings: On strong baselines, comparative analyses, and some misconceptions](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 710–721.
- Goran Glavas and Ivan Vulić. 2021. [Is supervised syntactic parsing beneficial for language understanding tasks? an empirical investigation](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, EACL 2021, Online, April 19 - 23, 2021*, pages 3090–3104. Association for Computational Linguistics.
- Dan Hendrycks and Kevin Gimpel. 2016. [Gaussian error linear units \(GELUs\)](#). *arXiv preprint*.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin de Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. [Parameter-efficient transfer learning for NLP](#). In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, pages 2790–2799.
- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020. [XTREME: A massively multilingual multi-task benchmark for evaluating cross-lingual generalization](#). In *Proceedings of the 37th International*

- Conference on Machine Learning, ICML 2020, 12-18 July 2020, Virtual Conference.*
- Eric Jang, Shixiang Gu, and Ben Poole. 2017. [Categorical reparameterization with gumbel-softmax](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*.
- Karthikeyan K, Zihan Wang, Stephen Mayhew, and Dan Roth. 2020. [Cross-lingual ability of multilingual BERT: an empirical study](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*.
- Anne Lauscher, Vinit Ravishankar, Ivan Vulić, and Goran Glavaš. 2020. [From zero to hero: On the limitations of zero-shot language transfer with multilingual Transformers](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4483–4499, Online.
- Benjamin Müller, Antonios Anastasopoulos, Benoît Sagot, and Djamé Seddah. 2021. [When being unseen from mbert is just the beginning: Handling new languages with multilingual language models](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021*, pages 448–462. Association for Computational Linguistics.
- Joakim Nivre, Marie-Catherine de Marneffe, Filip Ginter, Yoav Goldberg, Jan Hajič, Christopher D. Manning, Ryan McDonald, Slav Petrov, Sampo Pyysalo, Natalia Silveira, Reut Tsarfaty, and Daniel Zeman. 2016. [Universal Dependencies v1: A multilingual treebank collection](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 1659–1666, Portorož, Slovenia. European Language Resources Association (ELRA).
- Joakim Nivre, Marie-Catherine de Marneffe, Filip Ginter, Jan Hajič, Christopher D. Manning, Sampo Pyysalo, Sebastian Schuster, Francis Tyers, and Daniel Zeman. 2020. [Universal Dependencies v2: An evergrowing multilingual treebank collection](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 4034–4043, Marseille, France. European Language Resources Association.
- Xiaoman Pan, Boliang Zhang, Jonathan May, Joel Nothman, Kevin Knight, and Heng Ji. 2017. [Cross-lingual name tagging and linking for 282 languages](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pages 1946–1958.
- Jonas Pfeiffer, Aishwarya Kamath, Andreas Rücklé, Kyunghyun Cho, and Iryna Gurevych. 2021. [AdapterFusion: Non-Destructive Task Composition for Transfer Learning](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, EACL 2021, Online, April 19 - 23, 2021*, pages 487–503. Association for Computational Linguistics.
- Jonas Pfeiffer, Andreas Rücklé, Clifton Poth, Aishwarya Kamath, Ivan Vulic, Sebastian Ruder, Kyunghyun Cho, and Iryna Gurevych. 2020a. [AdapterHub: A Framework for Adapting Transformers](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, EMNLP 2020 - Demos, Online, November 16-20, 2020*, pages 46–54. Association for Computational Linguistics.
- Jonas Pfeiffer, Ivan Vulic, Iryna Gurevych, and Sebastian Ruder. 2020b. [MAD-X: An Adapter-Based Framework for Multi-Task Cross-Lingual Transfer](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 7654–7673. Association for Computational Linguistics.
- Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. [How multilingual is multilingual bert?](#) In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 4996–5001.
- Afshin Rahimi, Yuan Li, and Trevor Cohn. 2019. [Massively multilingual transfer for NER](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 151–164.
- Sylvestre-Alvise Rebuffi, Hakan Bilen, and Andrea Vedaldi. 2017. [Learning multiple visual domains with residual adapters](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 506–516.
- Sylvestre-Alvise Rebuffi, Hakan Bilen, and Andrea Vedaldi. 2018. [Efficient parametrization of multi-domain deep neural networks](#). In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 8119–8127.
- Sebastian Ruder, Ivan Vulić, and Anders Søgaard. 2019. [A survey of cross-lingual embedding models](#). *Journal of Artificial Intelligence Research*, 65:569–631.
- Phillip Rust, Jonas Pfeiffer, Ivan Vulic, Sebastian Ruder, and Iryna Gurevych. 2021. [How Good is Your Tokenizer? On the Monolingual Performance of Multilingual Language Models](#). In *Proceedings of the 59th Annual Meeting of the Association*

- for *Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 3118–3135. Association for Computational Linguistics.
- Anders Søgaard, Sebastian Ruder, and Ivan Vulic. 2018. [On the limitations of unsupervised bilingual dictionary induction](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 778–788.
- Asa Cooper Stickland and Iain Murray. 2019. [BERT and PALS: Projected Attention Layers for Efficient Adaptation in Multi-Task Learning](#). In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, pages 5986–5995.
- Ahmet Üstün, Arianna Bisazza, Gosse Bouma, and Gertjan van Noord. 2020. [UDapter: Language Adaptation for Truly Universal Dependency Parsing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 2302–2315. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention Is All You Need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 5998–6008.
- Marko Vidoni, Ivan Vulić, and Goran Glavaš. 2020. [Orthogonal language and task adapters in zero-shot cross-lingual transfer](#). In *arXiv preprint*.
- Zihan Wang, Karthikeyan K, Stephen Mayhew, and Dan Roth. 2020. [Extending multilingual BERT to low-resource languages](#). In *Findings of the 2020 Conference on Empirical Methods in Natural Language Processing*, pages 2649–2656.
- Shijie Wu, Alexis Conneau, Haoran Li, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Emerging cross-lingual structure in pretrained language models](#). In *Proceedings of the 58th Conference of the Association for Computational Linguistics, ACL 2020, Virtual Conference, July 6-8, 2020*, pages 6022–6034.
- Shijie Wu and Mark Dredze. 2019. [Beto, bentz, becas: The surprising cross-lingual effectiveness of BERT](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 833–844.
- Shijie Wu and Mark Dredze. 2020. [Are all languages created equal in multilingual BERT?](#) In *Proceedings of the 5th Workshop on Representation Learning for NLP*, pages 120–130, Online. Association for Computational Linguistics.
- Daniel Zeman, Joakim Nivre, Mitchell Abrams, Elia Ackermann, Noëmi Aeppli, Željko Agić, Lars Ahrenberg, Chika Kennedy Ajede, Gabrielè Aleksandravičiūtė, Lene Antonsen, Katya Aplonova, Angelina Aquino, Maria Jesus Aranzabe, Gashaw Arutie, Masayuki Asahara, Luma Ateyah, Furkan Atmaca, Mohammed Attia, Aitziber Atutxa, Liesbeth Augustinus, Elena Badmaeva, Miguel Ballesteros, Esha Banerjee, Sebastian Bank, Verginica Barbu Mititelu, Victoria Basmov, Colin Batchelor, John Bauer, Kepa Bengoetxea, Yevgeni Berzak, Irshad Ahmad Bhat, Riyaz Ahmad Bhat, Erica Biagetti, Eckhard Bick, Agnė Bielinškienė, Rogier Blokland, Victoria Bobicev, Loïc Boizou, Emanuel Borges Völker, Carl Börstell, Cristina Bosco, Gosse Bouma, Sam Bowman, Adriane Boyd, Kristina Brokaitė, Aljoscha Burchardt, Marie Candito, Bernard Caron, Gauthier Caron, Tatiana Cavalcanti, Gülşen Cebiroğlu Eryiğit, Flavio Massimiliano Cecchini, Giuseppe G. A. Celano, Slavomír Čepľo, Savas Cetin, Fabricio Chalub, Ethan Chi, Jinho Choi, Yongseok Cho, Jayeol Chun, Alessandra T. Cignarella, Silvie Cinková, Aurélie Collomb, Çağrı Çöltekin, Miriam Connor, Marine Courtin, Elizabeth Davidson, Marie-Catherine de Marneffe, Valeria de Paiva, Elvis de Souza, Arantza Diaz de Ilaraza, Carly Dickerson, Bamba Dione, Peter Dirix, Kaja Dobrovoljc, Timothy Dozat, Kira Drokanova, Puneet Dwivedi, Hanne Eckhoff, Marhaba Eli, Ali Elkahky, Binyam Ephrem, Olga Erina, Tomaž Erjavec, Aline Etienne, Wograinne Evelyn, Richárd Farkas, Hector Fernandez Alcalde, Jennifer Foster, Cláudia Freitas, Kazunori Fujita, Katarína Gajdošová, Daniel Galbraith, Marcos Garcia, Moa Gärdenfors, Sebastian Garza, Kim Gerdes, Filip Ginter, Iakes Goenaga, Koldo Gojenola, Memduh Gökırmak, Yoav Goldberg, Xavier Gómez Guinovart, Berta González Saavedra, Bernadeta Griciūtė, Matias Grioni, Loïc Grobol, Normunds Grūzītis, Bruno Guillaume, Céline Guillot-Barbance, Tunga Güngör, Nizar Habash, Jan Hajič, Jan Hajič jr., Mika Hämmäläinen, Linh Hà Mỹ, Na-Rae Han, Kim Harris, Dag Haug, Johannes Heinecke, Oliver Hellwig, Felix Hennig, Barbora Hladká, Jaroslava Hlaváčová, Florinel Hociung, Petter Hohle, Jena Hwang, Takumi Ikeda, Radu Ion, Elena Irimia, Olájidé Ishola, Tomáš Jelínek, Anders Johannsen, Hildur Jónsdóttir, Fredrik Jørgensen, Markus Juutinen, Hüner Kaşıkara, Andre Kaasen, Nadezhda Kabaeva, Sylvain Kahane, Hiroshi Kanayama, Jenna Kanerva, Boris Katz, Tolga Kayadelen, Jessica Kenney, Václava Kettnerová, Jesse Kirchner, Elena Klementieva, Arne Köhn, Abdullatif Köksal, Kamil Kopacewicz, Timo Korkiakangas, Natalia Kotsyba, Jolanta Kovalevskaitė, Simon Krek, Sookyong Kwak, Veronika Laippala, Lorenzo Lambertino, Lucia Lam, Tatiana Lando, Septina Dian Larasati, Alexei Lavrentiev, John Lee, Phuong Lê Hồng, Alessandro Lenci, Saran Lertpradit, Herman Leung, Maria Levina, Cheuk Ying Li, Josie Li, Keying Li, KyungTae Lim, Yuan Li, Nikola

Ljubešić, Olga Loginova, Olga Lyashevskaya, Teresa Lynn, Vivien Macketanz, Aibek Makazhanov, Michael Mandl, Christopher Manning, Ruli Manurung, Cătălina Mărănduc, David Mareček, Katrin Marheinecke, Héctor Martínez Alonso, André Martins, Jan Mašek, Hiroshi Matsuda, Yuji Matsumoto, Ryan McDonald, Sarah McGuinness, Gustavo Mendonça, Niko Miekka, Margarita Misirpashayeva, Anna Missilä, Cătălin Mititelu, Maria Mitrofan, Yusuke Miyao, Simonetta Montemagni, Amir More, Laura Moreno Romero, Keiko Sophie Mori, Tomohiko Morioka, Shinsuke Mori, Shigeki Moro, Bjartur Mortensen, Bohdan Moskalevskyi, Kadri Muischnek, Robert Munro, Yugo Murawaki, Kaili Müürisep, Pinkey Nainwani, Juan Ignacio Navarro Horňiáček, Anna Nedoluzhko, Gunta Nešpore-Bērzkalne, Luong Nguyễn Thị, Huyèn Nguyễn Thị Minh, Yoshihiro Nikaido, Vitaly Nikolaev, Rattima Nitisaroj, Hanna Nurmi, Stina Ojala, Atul Kr. Ojha, Adédayo Olúòkun, Mai Omura, Emeka Onwuegbuzia, Petya Osenova, Robert Östling, Lilja Övrelid, Şaziye Betül Özateş, Arzucan Özgür, Balkız Öztürk Başaran, Niko Partanen, Elena Pascual, Marco Passarotti, Agnieszka Patejuk, Guilherme Paulino-Passos, Angelika Peljak-Łapińska, Siyao Peng, Cene-Augusto Perez, Guy Perrier, Daria Petrova, Slav Petrov, Jason Phelan, Jussi Pitulainen, Tommi A Pirinen, Emily Pitler, Barbara Plank, Thierry Poibeau, Larisa Ponomareva, Martin Popel, Lauma Pretkalniņa, Sophie Prévost, Prokopis Prokopidis, Adam Przepiórkowski, Tiina Puolakainen, Sampo Pyysalo, Peng Qi, Andriela Rääbis, Alexandre Rademaker, Loganathan Ramasamy, Taraka Rama, Carlos Ramisch, Vinit Ravishankar, Livy Real, Petru Rebeja, Siva Reddy, Georg Rehm, Ivan Riabov, Michael Riebler, Erika Rimkutė, Larissa Rinaldi, Laura Rituma, Luisa Rocha, Mykhailo Romanenko, Rudolf Rosa, Valentin Rosca, Davide Rovati, Olga Rudina, Jack Rueter, Shoal Sadde, Benoît Sagot, Shadi Saleh, Alessio Salomoni, Tanja Samardžić, Stephanie Samson, Manuela Sanguinetti, Dage Särg, Baiba Saulīte, Yanin Sawanakunanon, Salvatore Scarlata, Nathan Schneider, Sebastian Schuster, Djamé Seddah, Wolfgang Seeker, Mojgan Seraji, Mo Shen, Atsuko Shimada, Hiroyuki Shirasu, Muh Shohibus-sirri, Dmitry Sichinava, Aline Silveira, Natalia Silveira, Maria Simi, Radu Simionescu, Katalin Simkó, Mária Šimková, Kiril Simov, Maria Skachedubova, Aaron Smith, Isabela Soares-Bastos, Carolyn Spadine, Antonio Stella, Milan Straka, Jana Strnadová, Alane Suhr, Umut Sulubacak, Shingo Suzuki, Zsolt Szántó, Dima Taji, Yuta Takahashi, Fabio Tamburini, Takaaki Tanaka, Samson Tella, Isabelle Tellier, Guillaume Thomas, Liisi Torga, Marsida Toska, Trond Trosterud, Anna Trukhina, Reut Tsarfaty, Utku Türk, Francis Tyers, Sumire Uematsu, Roman Untilov, Zdeňka Urešová, Larraitx Uriá, Hans Uszkoreit, Andrius Utka, Sowmya Vajjala, Daniel van Niekerk, Gertjan van Noord, Viktor Varga, Eric Villemonte de la Clergerie, Veronika Vincze, Aya Wakasa, Lars Wallin, Abigail Walsh, Jing Xian Wang, Jonathan North Washington, Max-

imilan Wendt, Paul Widmer, Seyi Williams, Mats Wirén, Christian Wittern, Tsegay Woldemariam, Tak-sum Wong, Alina Wróblewska, Mary Yako, Kayo Yamashita, Naoki Yamazaki, Chunxiao Yan, Koichi Yasuoka, Marat M. Yavrumyan, Zhuoran Yu, Zdeněk Žabokrtský, Amir Zeldes, Hanzhi Zhu, and Anna Zhuravleva. 2020. *Universal dependencies 2.6*. LINDAT/CLARIAH-CZ digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.

## A Appendices

### A.1 Full Model Transfer

For comparability with previous work, which generally fine-tunes the entire model (Chau et al., 2020; Wang et al., 2020), we follow Artetxe et al. (2020) and learn a new embedding matrix for the target language while freezing the pretrained transformer weights. For training on the target task, we fine-tune the transformer weights of the pretrained model while keeping the original embedding layer frozen. For zero-shot transfer, we replace the existing with the new embedding matrix trained on the target language.

### A.2 Results: Named Entity Recognition

We present non-aggregated NER transfer performance when transferring from English, Chinese, Japanese, and Arabic in Tables 7, 8, 9, and 10 respectively. *12Ad* indicates whether (✓) or not (✗) an adapter is placed in the 12th transformer layer. We additionally present the results for full model fine-tuning (FMT-\*).

### A.3 Results: Dependency Parsing

We present non-aggregated DP transfer performance when transferring from English, Chinese, Japanese, and Arabic in Tables 11a, 11b, 11c, and 11d respectively.

### A.4 Script Clusters

We present the groups of scripts within the 10 KMeans clusters in Figure 4. We follow Ács (2019) in grouping the scripts into languages.

### A.5 Language information

Table 12 lists all 104 languages and corresponding scripts which mBERT was pretrained on.

### A.6 Hardware Setup

All experiments were conducted on a single NVIDIA V100 GPU with 32 Gb of VRAM.

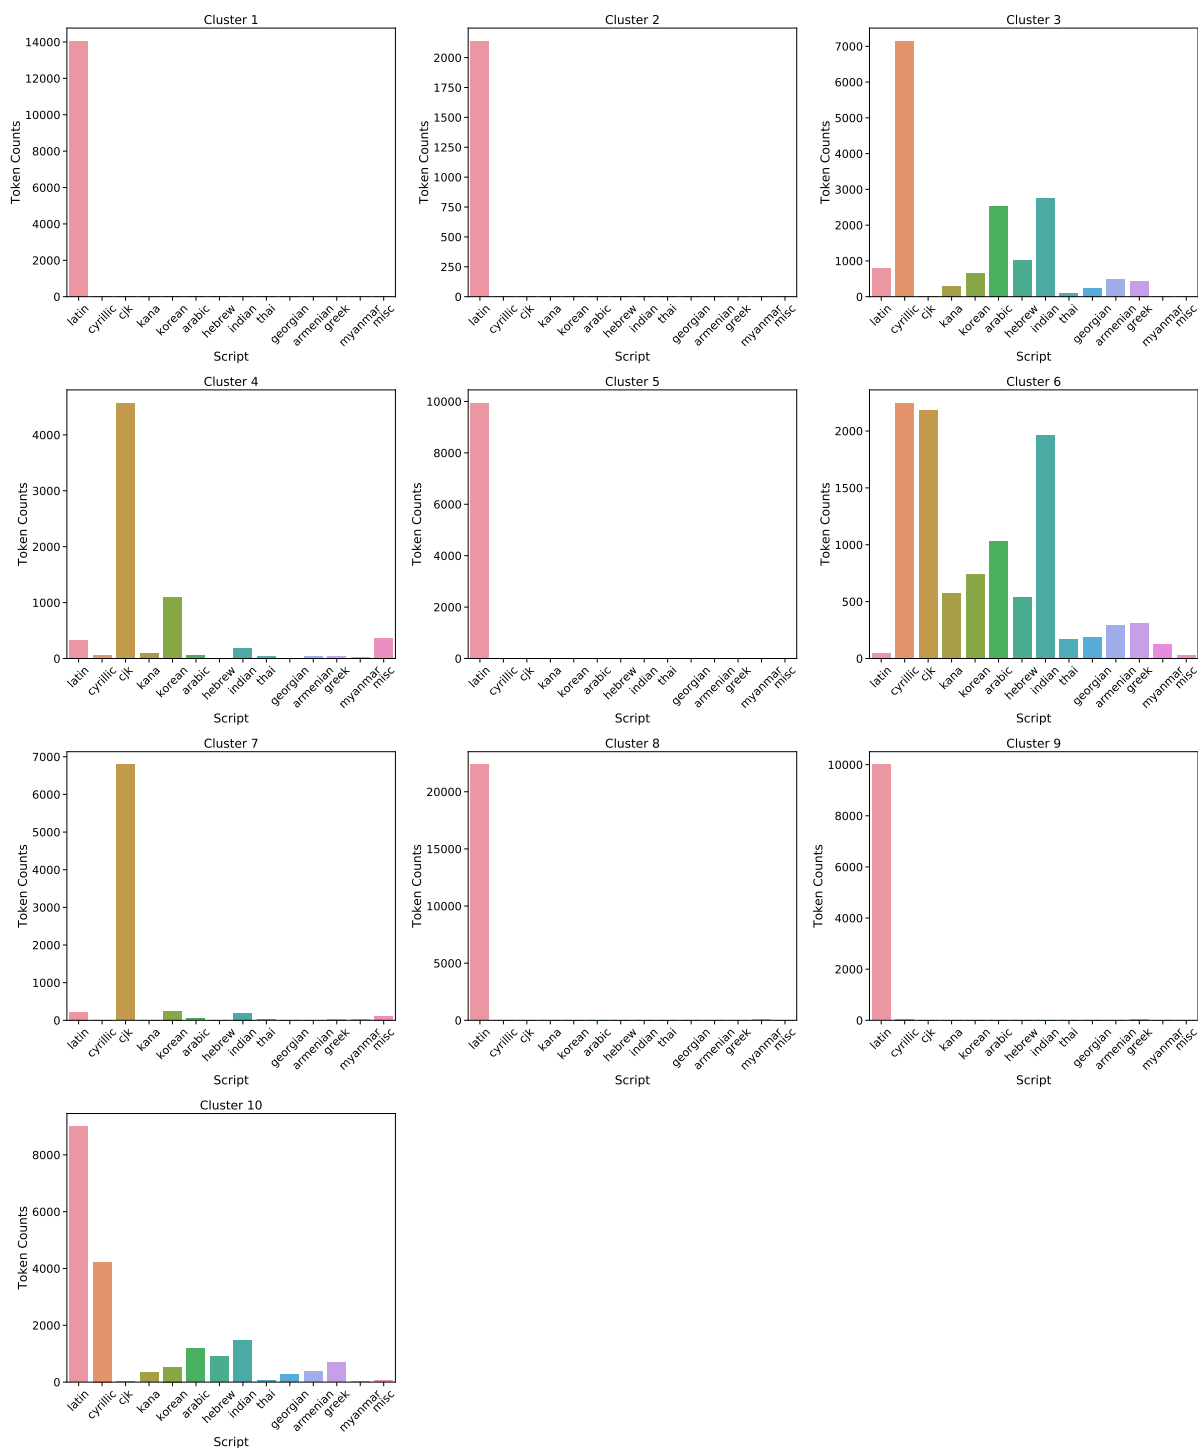


Figure 4: Number of tokens of each script grouped into the 10 KMeans clusters. We follow (Ács, 2019) in splitting the utf-8 tokens into scripts.

	12Ad	Seen Languages				Unseen Languages but Covered Scripts										New Scripts					
		ka	ur	hi	Avg	cdo	mi	ilo	gn	xmf	sd	myv	bh	wo	Avg	am	bo	km	dv	si	Avg
mBERT		64.3	34.3	65.3	54.6	17.3	22.2	62.7	46.7	34.4	12.1	47.2	48.8	26.7	35.4	0.0	11.4	4.9	0.3	0.3	3.4
MAD-X	✓	64.8	44.5	60.7	56.7	25.1	38.4	61.2	52.7	46.0	25.3	47.7	52.1	47.1	43.9	14.5	20.4	14.5	13.9	16.7	16.0
MAD-X 2.0	✗	67.4	54.4	64.3	62.1	24.5	51.8	79.2	54.4	49.0	30.9	58.1	55.6	56.8	51.1	8.0	16.6	18.4	23.0	22.0	17.6
FMT-RANDINIT		65.1	48.9	63.3	59.1	26.9	12.6	67.1	20.9	54.6	37.5	8.0	47.1	10.9	31.7	28.6	29.3	51.6	27.0	50.0	37.3
EL-RAND	✓	66.6	35.0	62.9	54.8	33.5	11.0	58.8	49.0	47.8	37.5	20.0	53.6	4.1	35.0	36.8	40.2	57.6	35.5	48.4	43.7
EL-RAND	✗	70.1	39.2	67.5	58.9	34.0	16.4	66.6	53.0	57.9	39.7	35.1	56.4	4.9	40.4	46.4	50.8	61.4	38.9	52.6	50.0
MF <sup>1</sup> -RAND	✓	62.1	46.2	61.2	56.5	29.5	28.1	51.4	41.6	47.3	40.8	38.5	33.8	29.0	37.8	39.6	39.3	54.7	35.4	46.6	43.1
MF <sup>1</sup> -RAND	✗	67.2	51.5	63.2	60.6	33.3	33.7	58.0	48.6	56.1	44.3	55.2	43.2	41.8	46.0	43.3	42.1	63.4	41.6	54.7	49.0
MF <sup>10</sup> <sub>KMEANS</sub> -RAND	✓	63.2	52.0	61.4	58.9	28.4	7.2	52.8	45.6	51.3	15.3	30.0	47.8	32.8	34.6	37.4	34.1	55.1	33.7	44.4	40.9
MF <sup>10</sup> <sub>KMEANS</sub> -RAND	✗	67.3	55.5	63.9	62.3	35.1	6.3	65.0	57.3	59.1	20.4	44.8	55.2	42.6	42.9	46.8	41.8	60.1	37.5	51.0	47.4
MF <sup>10</sup> <sub>NEURAL</sub> -RAND	✓	63.9	51.7	62.6	59.4	14.0	4.1	57.0	49.4	48.2	35.6	23.6	52.4	34.8	35.5	31.6	32.0	58.3	17.2	35.6	34.9
MF <sup>10</sup> <sub>NEURAL</sub> -RAND	✗	68.4	56.3	62.8	62.5	16.0	4.7	65.9	56.4	57.6	42.2	39.1	58.5	44.8	42.8	37.3	31.0	61.7	18.9	42.2	38.2
FMT-LEXINIT		67.6	51.8	63.1	60.9	42.2	31.3	72.4	47.6	56.7	35.6	53.1	53.4	45.6	48.7	34.9	25.9	49.9	43.9	55.4	42.0
EL-LEX	✓	70.1	51.4	65.2	62.2	42.5	37.9	56.8	50.8	52.5	41.1	38.8	56.1	54.2	47.9	40.0	41.9	60.8	43.3	52.6	47.7
EL-LEX	✗	73.1	56.0	69.2	66.1	42.9	34.3	66.1	54.9	56.9	44.0	50.8	59.4	61.4	52.3	46.5	46.6	60.0	48.3	60.3	52.3
MF <sup>1</sup> -LEX	✓	66.4	41.1	62.0	56.5	41.3	37.2	57.1	46.5	47.0	35.7	29.5	50.6	36.5	42.4	36.4	40.0	59.1	35.6	44.1	43.0
MF <sup>1</sup> -LEX	✗	71.0	51.7	65.6	62.8	43.3	55.6	64.9	57.5	53.2	42.1	47.5	57.4	45.9	51.9	43.1	41.8	63.1	44.0	53.6	49.1
MF <sup>10</sup> <sub>KMEANS</sub> -LEX	✓	68.0	56.4	63.7	62.7	38.1	45.5	60.2	47.8	46.5	39.8	33.5	52.2	50.1	46.0	42.2	36.8	61.6	39.1	50.5	46.0
MF <sup>10</sup> <sub>KMEANS</sub> -LEX	✗	71.0	58.0	65.8	64.9	44.7	65.7	67.1	55.5	55.2	43.9	41.0	57.3	59.2	54.4	49.0	36.3	65.3	45.5	55.3	50.3
MF <sup>10</sup> <sub>NEURAL</sub> -LEX	✓	65.7	53.6	60.2	59.8	43.7	51.1	52.7	51.0	46.7	41.1	46.6	55.9	48.8	48.6	43.8	37.3	58.9	39.1	41.1	44.0
MF <sup>10</sup> <sub>NEURAL</sub> -LEX	✗	70.6	60.5	63.9	65.0	45.3	65.0	61.9	53.6	53.0	43.2	56.9	60.2	60.2	55.5	53.5	46.1	61.1	44.5	43.2	49.7

Table 7: Mean  $F_1$  NER test results averaged over 5 runs transferring from high resource language **English** to the low resource languages. *12Ad* indicates whether (✓) or not (✗) an adapter is placed in the 12th transformer layer. The top group (first three rows) includes models which leverage the original tokenizer which is not specialized for the target language. The second group (last 18 rows) include models with new tokenizers. Here we separate models with randomly initialized embeddings (\*-RANDINIT) from models with lexical initialization (\*-LEXINIT) by the dashed line. We additionally present the results for full model fine-tuning (FMT-\*).

	12Ad	Seen Languages				Unseen Languages but Covered Scripts										New Scripts					
		ka	ur	hi	Avg	cdo	mi	ilo	gn	xmf	sd	myv	bh	wo	Avg	am	bo	km	dv	si	Avg
mBERT		66.0	33.3	60.2	53.2	21.2	42.3	42.6	59.5	49.1	14.2	29.4	60.1	25.1	38.2	0.2	13.1	21.8	1.4	2.0	7.7
MAD-X	✓	62.9	50.8	54.1	55.9	32.2	42.5	40.6	56.4	57.3	44.4	68.6	63.1	55.1	51.1	14.6	24.4	24.1	19.8	18.9	20.4
MAD-X 2.0	✗	63.9	52.4	55.7	57.3	37.0	42.1	64.7	57.9	54.9	42.8	72.0	63.8	60.1	55.1	13.5	29.0	23.0	21.5	20.2	21.4
FMT-RANDINIT		64.6	37.2	56.5	52.8	36.7	4.1	48.1	15.7	57.8	45.8	7.0	48.9	15.8	31.1	26.4	28.8	53.0	31.1	41.6	36.2
EL-RAND	✓	62.8	43.7	57.2	54.6	41.4	6.8	51.7	45.5	57.4	48.1	38.7	58.4	13.4	40.2	41.8	49.8	51.6	30.1	54.3	45.5
EL-RAND	✗	63.0	48.3	59.3	56.9	44.8	4.4	54.3	47.5	60.1	47.0	46.5	58.6	8.7	41.3	43.5	51.8	53.6	29.7	51.3	46.0
MF <sup>1</sup> -RAND	✓	59.3	38.1	50.3	49.2	38.2	10.0	34.7	44.4	54.1	46.5	59.8	43.0	47.8	42.1	41.6	49.1	59.7	42.3	53.6	49.2
MF <sup>1</sup> -RAND	✗	59.5	42.8	52.9	51.7	44.2	10.2	35.9	44.3	54.2	44.8	64.2	43.4	47.4	43.2	43.9	48.6	55.9	45.3	52.4	49.2
MF <sup>10</sup> <sub>KMEANS</sub> -RAND	✓	60.0	40.7	55.5	52.1	38.6	1.0	43.2	44.4	64.3	28.6	55.0	61.0	50.9	43.0	39.3	23.5	63.7	34.6	45.0	41.2
MF <sup>10</sup> <sub>KMEANS</sub> -RAND	✗	59.6	43.9	56.8	53.4	48.1	3.1	47.6	50.3	63.1	33.0	60.6	59.7	51.9	46.4	44.9	25.3	62.1	37.5	45.7	43.1
MF <sup>10</sup> <sub>NEURAL</sub> -RAND	✓	62.5	44.1	54.1	53.5	21.3	2.2	44.7	46.1	58.4	50.8	39.7	63.9	47.6	41.6	37.5	38.7	65.4	18.3	50.8	42.1
MF <sup>10</sup> <sub>NEURAL</sub> -RAND	✗	62.1	50.0	56.3	56.1	20.7	4.1	49.7	51.4	57.9	52.5	48.2	65.6	48.3	44.3	38.3	42.9	64.3	18.1	52.5	43.2
FMT-LEXINIT		67.1	36.9	60.2	54.7	59.6	15.6	49.3	55.4	57.3	45.5	52.3	64.6	50.4	50.0	29.3	31.3	58.2	45.2	48.6	42.5
EL-LEX	✓	68.9	49.7	61.7	60.1	51.4	26.6	54.3	53.5	54.6	48.4	59.7	62.2	59.3	52.2	43.9	47.1	60.2	48.3	55.7	51.0
EL-LEX	✗	68.0	53.1	64.0	61.7	54.7	24.4	54.4	55.3	52.4	48.4	60.7	59.6	58.1	52.0	48.0	53.5	60.7	50.3	54.8	53.5
MF <sup>1</sup> -LEX	✓	61.8	42.8	56.4	53.7	51.2	16.3	47.2	49.7	52.6	47.2	56.2	60.0	53.7	48.2	41.5	39.4	61.4	41.7	44.8	45.8
MF <sup>1</sup> -LEX	✗	61.9	47.8	57.7	55.8	48.2	18.8	51.5	55.8	54.9	46.7	61.3	60.4	52.6	50.0	42.9	41.9	58.3	44.8	49.6	47.5
MF <sup>10</sup> <sub>KMEANS</sub> -LEX	✓	62.9	47.2	58.3	56.1	49.2	31.8	53.8	53.2	51.1	50.1	56.0	61.6	58.5	51.7	42.7	43.6	62.4	46.9	50.0	49.1
MF <sup>10</sup> <sub>KMEANS</sub> -LEX	✗	63.3	50.9	59.8	58.0	51.8	29.6	52.5	55.3	51.5	50.7	59.9	65.1	56.0	52.5	44.9	44.6	61.0	49.9	48.0	49.6
MF <sup>10</sup> <sub>NEURAL</sub> -LEX	✓	62.6	55.1	55.6	57.8	52.9	32.3	40.0	54.4	59.4	53.4	61.1	60.8	67.4	53.5	44.3	50.1	65.7	40.9	53.4	50.9
MF <sup>10</sup> <sub>NEURAL</sub> -LEX	✗	63.7	57.1	58.4	59.7	54.1	34.2	42.0	54.6	59.0	49.6	63.4	62.9	66.9	54.1	41.5	53.8	62.3	45.4	49.6	50.5

Table 8: Mean  $F_1$  NER test results averaged over 5 runs transferring from high resource language **Chinese** to the low resource languages. *12Ad* indicates whether (✓) or not (✗) an adapter is placed in the 12th transformer layer. The top group (first three rows) includes models which leverage the original tokenizer which is not specialized for the target language. The second group (last 18 rows) include models with new tokenizers. Here we separate models with randomly initialized embeddings (\*-RANDINIT) from models with lexical initialization (\*-LEXINIT) by the dashed line. We additionally present the results for full model fine-tuning (FMT-\*).

	12Ad	Seen Languages				Unseen Languages but Covered Scripts											New Scripts				
		ka	ur	hi	Avg	cdo	mi	ilo	gn	xmf	sd	myy	bh	wo	Avg	am	bo	km	dv	si	Avg
mBERT		64.2	35.5	59.8	53.2	17.7	42.6	39.5	50.2	46.5	14.2	23.9	56.0	21.6	34.7	3.4	23.2	6.2	1.4	7.6	8.3
MAD-X	✓	62.7	49.8	55.9	56.1	29.7	43.8	28.3	48.0	53.2	34.2	58.0	54.7	52.9	44.7	6.0	31.4	22.4	16.7	17.2	18.7
MAD-X 2.0	✗	64.7	52.5	57.8	58.3	31.8	46.6	43.5	54.6	51.7	40.0	60.6	59.9	50.8	48.8	7.2	31.2	22.7	24.4	22.1	21.5
FMT-RANDINIT		64.2	36.4	60.2	53.6	33.3	3.8	49.7	15.5	56.1	47.2	11.7	54.6	12.7	31.6	27.5	33.0	58.5	30.4	52.5	40.4
EL-RAND	✓	63.0	37.1	58.2	52.8	35.8	3.9	47.3	43.7	50.1	41.6	22.5	54.7	5.7	33.9	34.9	50.6	52.7	32.3	44.6	43.0
EL-RAND	✗	65.6	43.1	60.7	56.5	37.4	2.6	52.6	47.4	58.1	46.7	31.2	56.8	4.3	37.5	40.3	53.5	60.7	32.3	52.8	47.9
MF <sup>1</sup> -RAND	✓	60.8	43.1	52.7	52.2	36.7	9.1	32.4	42.0	55.2	42.1	57.1	43.1	46.7	40.5	37.2	48.3	56.8	37.9	50.6	46.2
MF <sup>1</sup> -RAND	✗	63.5	44.6	55.7	54.6	43.3	12.2	28.9	44.9	57.7	48.5	62.4	46.0	55.9	44.4	41.4	46.8	57.6	40.1	52.8	47.7
MF <sup>10</sup> <sub>KMEANS</sub> -RAND	✓	58.8	36.5	56.5	50.6	34.5	0.9	24.0	45.7	57.8	20.6	27.3	52.4	42.2	33.9	33.5	39.9	61.9	34.7	40.7	42.1
MF <sup>10</sup> <sub>KMEANS</sub> -RAND	✗	64.4	39.0	58.4	53.9	40.3	1.3	31.3	48.8	57.0	26.5	46.1	55.5	43.8	39.0	36.2	40.7	63.6	37.0	47.2	44.9
MF <sup>10</sup> <sub>NEURAL</sub> -RAND	✓	59.4	37.4	55.4	50.7	22.6	2.6	29.3	44.8	52.1	45.5	28.3	58.5	41.3	36.1	31.0	42.4	63.9	17.0	45.5	40.0
MF <sup>10</sup> <sub>NEURAL</sub> -RAND	✗	62.8	44.4	56.7	54.6	19.5	3.7	32.6	46.0	57.3	44.5	45.4	62.3	42.9	39.4	33.8	40.1	63.7	18.2	44.5	40.0
FMT-LEXINIT		66.1	40.4	62.0	56.2	52.4	13.7	45.5	52.8	56.7	43.0	61.2	64.0	50.2	48.9	31.5	30.0	60.7	40.2	53.9	43.3
EL-LEX	✓	60.3	51.7	61.8	57.9	42.3	20.1	55.3	47.5	50.8	42.7	60.9	59.8	55.2	48.3	39.8	44.8	60.2	39.9	52.7	47.5
EL-LEX	✗	68.8	52.5	62.9	61.4	46.6	22.6	51.3	50.4	55.0	44.7	60.1	63.4	56.0	50.0	45.0	47.3	61.7	45.4	57.7	51.4
MF <sup>1</sup> -LEX	✓	62.7	47.4	58.6	56.2	47.0	22.1	38.0	49.6	51.6	42.0	57.1	59.2	55.2	46.9	34.9	41.4	63.3	41.1	43.0	44.8
MF <sup>1</sup> -LEX	✗	64.6	52.1	59.4	58.7	48.8	26.7	43.0	51.0	55.1	44.6	58.3	63.1	57.3	49.8	32.2	42.0	63.8	43.1	50.4	46.3
MF <sup>10</sup> <sub>KMEANS</sub> -LEX	✓	62.2	46.9	57.4	55.5	43.2	33.6	31.8	46.1	53.5	42.3	35.8	60.2	55.8	44.7	37.8	41.9	60.9	44.7	50.5	47.2
MF <sup>10</sup> <sub>KMEANS</sub> -LEX	✗	67.0	51.5	59.4	59.3	51.2	45.8	35.6	49.9	55.0	45.8	50.5	64.3	56.7	50.5	42.6	38.7	63.6	44.3	51.2	48.1
MF <sup>10</sup> <sub>NEURAL</sub> -LEX	✓	62.1	49.2	56.4	55.9	45.1	20.5	29.2	47.1	51.3	45.9	48.2	62.3	62.2	45.8	33.6	50.4	67.1	40.0	45.9	47.4
MF <sup>10</sup> <sub>NEURAL</sub> -LEX	✗	66.2	52.4	58.5	59.0	48.4	36.0	29.6	51.5	54.9	46.8	56.3	63.6	65.9	50.3	38.1	54.7	64.6	43.6	46.8	49.6

Table 9: Mean  $F_1$  NER test results averaged over 5 runs transferring from high resource language **Japanese** to the low resource languages. *12Ad* indicates whether (✓) or not (✗) an adapter is placed in the 12th transformer layer. The top group (first three rows) includes models which leverage the original tokenizer which is not specialized for the target language. The second group (last 18 rows) include models with new tokenizers. Here we separate models with randomly initialized embeddings (\*-RANDINIT) from models with lexical initialization (\*-LEXINIT) by the dashed line. We additionally present the results for full model fine-tuning (FMT-\*).

	12Ad	Seen Languages				Unseen Languages but Covered Scripts											New Scripts				
		ka	ur	hi	Avg	cdo	mi	ilo	gn	xmf	sd	myy	bh	wo	Avg	am	bo	km	dv	si	Avg
mBERT		64.4	46.2	70.0	60.2	17.0	25.8	36.7	58.4	43.6	10.3	22.3	53.2	23.6	32.3	0.0	21.9	9.8	1.9	0.3	6.8
MAD-X	✓	65.3	60.6	64.1	63.4	29.5	57.8	52.0	53.5	51.2	32.3	57.0	57.9	43.2	48.3	7.9	22.9	9.6	16.8	14.4	14.3
MAD-X 2.0	✗	66.5	61.9	66.1	64.9	29.5	61.8	68.7	58.4	54.8	34.8	61.3	59.8	54.5	53.7	14.1	22.0	8.5	20.0	10.4	15.0
FMT-RANDINIT		64.3	44.0	66.9	58.4	37.2	2.0	59.6	20.0	64.9	41.1	9.1	55.8	4.9	32.7	30.2	38.3	56.7	31.7	54.7	42.3
EL-RAND	✓	61.8	59.1	64.9	61.9	31.8	6.7	52.7	43.3	55.4	38.0	29.4	48.6	3.0	34.3	36.9	59.8	61.4	27.7	44.3	46.0
EL-RAND	✗	64.4	60.7	67.8	64.3	35.9	5.9	58.1	45.9	61.5	43.6	27.8	50.4	2.8	36.9	41.4	59.4	62.2	30.1	49.2	48.5
MF <sup>1</sup> -RAND	✓	61.8	52.9	57.8	57.5	36.4	18.5	46.8	49.3	59.0	39.0	57.5	40.3	34.2	42.3	36.4	58.2	66.7	45.2	47.0	50.7
MF <sup>1</sup> -RAND	✗	63.8	55.1	59.5	59.5	38.4	21.4	51.8	51.6	63.5	42.8	59.8	44.9	35.0	45.5	42.0	58.4	69.8	46.7	48.6	53.1
MF <sup>10</sup> <sub>KMEANS</sub> -RAND	✓	60.4	55.5	61.5	59.1	32.3	10.5	43.5	47.6	60.0	25.5	29.7	44.2	21.0	34.9	36.6	46.5	68.3	40.1	46.3	47.5
MF <sup>10</sup> <sub>KMEANS</sub> -RAND	✗	61.3	59.4	63.5	61.4	39.3	11.4	49.9	51.4	62.4	28.8	35.3	47.8	26.6	39.2	43.1	45.7	69.1	42.7	48.2	49.8
MF <sup>10</sup> <sub>NEURAL</sub> -RAND	✓	62.9	56.4	63.2	60.8	24.5	65.0	49.0	46.1	58.1	40.7	56.9	49.0	29.5	46.6	35.0	35.0	65.5	15.5	43.2	38.8
MF <sup>10</sup> <sub>NEURAL</sub> -RAND	✗	64.1	59.6	65.8	63.2	26.8	34.2	56.6	49.2	59.8	41.5	63.4	53.7	33.8	46.6	39.6	36.1	67.1	18.9	49.6	42.3
FMT-LEXINIT		64.3	47.6	67.5	59.8	47.0	28.5	68.9	53.4	66.0	43.3	49.1	57.5	43.0	50.7	34.6	32.1	62.6	45.9	50.4	45.2
EL-LEX	✓	62.0	68.0	67.1	65.7	40.7	47.5	52.1	54.3	58.4	35.8	53.8	53.5	24.8	46.8	46.1	59.2	61.5	42.1	52.4	52.3
EL-LEX	✗	67.5	70.0	69.5	69.0	42.9	47.7	63.4	55.7	63.6	41.4	53.6	55.4	25.1	49.9	46.5	58.6	61.8	44.8	52.3	52.8
MF <sup>1</sup> -LEX	✓	63.3	63.4	61.7	62.8	41.0	58.9	53.6	50.1	64.7	40.3	54.8	53.9	31.3	49.8	33.5	38.9	67.8	42.4	47.3	46.0
MF <sup>1</sup> -LEX	✗	64.7	64.3	63.4	64.1	46.3	58.1	61.3	54.6	65.7	46.4	53.1	55.5	36.5	53.1	36.1	44.4	70.2	43.1	52.2	49.2
MF <sup>10</sup> <sub>KMEANS</sub> -LEX	✓	64.8	61.9	64.2	63.7	35.9	52.0	56.0	51.7	57.5	41.5	50.3	51.1	30.3	47.4	46.2	52.6	69.2	44.1	49.2	52.3
MF <sup>10</sup> <sub>KMEANS</sub> -LEX	✗	66.5	62.9	65.4	64.9	43.3	62.0	57.5	55.0	60.7	45.2	50.1	55.9	35.1	51.6	50.5	52.4	70.9	45.5	51.6	54.2
MF <sup>10</sup> <sub>NEURAL</sub> -LEX	✓	63.4	66.5	63.4	64.4	44.7	38.6	40.7	52.0	58.7	42.5	52.9	54.6	34.8	46.6	46.4	53.4	64.3	44.4	42.5	50.2
MF <sup>10</sup> <sub>NEURAL</sub> -LEX	✗	64.7	68.3	63.7	65.5	47.9	54.9	52.0	54.6	59.8	45.3	57.0	57.3	38.7	51.9	48.9	51.9	68.6	46.2	45.3	52.2

Table 10: Mean  $F_1$  NER test results averaged over 5 runs transferring from high resource language **Arabic** to the low resource languages. *12Ad* indicates whether (✓) or not (✗) an adapter is placed in the 12th transformer layer. The top group (first three rows) includes models which leverage the original tokenizer which is not specialized for the target language. The second group (last 18 rows) include models with new tokenizers. Here we separate models with randomly initialized embeddings (\*-RANDINIT) from models with lexical initialization (\*-LEXINIT) by the dashed line. We additionally present the results for full model fine-tuning (FMT-\*).



	Seen Languages			Unseen Languages but Covered Scripts				New Script am
	hi	ur	Avg	bh	myv	wo	Avg	
	UAS / LAS	UAS / LAS	UAS / LAS	UAS / LAS	UAS / LAS	UAS / LAS	UAS / LAS	
mBERT	48.0 / 34.4	36.4 / 23.9	42.2 / 29.2	33.0 / 18.3	33.6 / 14.3	33.0 / 9.3	33.2 / 14.0	15.5 / 12.6
MAD-X 2.0	44.6 / 31.2	37.6 / 23.9	41.1 / 27.5	29.5 / 16.1	60.5 / 43.2	50.5 / 33.8	46.8 / 31.1	11.7 / 9.6
EL-RAND	45.9 / 32.5	38.2 / 23.9	42.1 / 28.2	29.9 / 13.6	57.8 / 35.0	33.2 / 11.1	40.3 / 19.9	26.4 / 13.0
MF <sup>1</sup> -RAND	45.3 / 33.3	37.8 / 25.2	41.6 / 29.3	26.6 / 13.3	62.7 / 41.5	42.2 / 27.2	43.8 / 27.3	32.7 / 15.5
MF <sup>10</sup> <sub>KMEANS</sub> -RAND	46.3 / 33.1	37.8 / 25.6	42.1 / 29.3	31.4 / 16.0	62.9 / 42.8	24.9 / 13.7	39.7 / 24.1	29.7 / 14.9
MF <sup>10</sup> <sub>NEURAL</sub> -RAND	42.6 / 30.5	33.7 / 22.6	38.1 / 26.5	33.2 / 17.8	61.6 / 41.5	38.3 / 22.2	44.4 / 27.2	28.4 / 12.7
EL-LEX	47.1 / 33.1	38.1 / 24.9	42.6 / 29.0	31.3 / 17.2	63.3 / 40.9	51.3 / 34.0	48.6 / 30.7	32.4 / 9.1
MF <sup>1</sup> -LEX	45.8 / 33.5	38.1 / 26.2	42.0 / 29.8	33.2 / 17.6	62.6 / 42.4	50.0 / 35.2	48.6 / 31.8	35.1 / 16.1
MF <sup>10</sup> <sub>KMEANS</sub> -LEX	46.0 / 33.4	38.8 / 26.5	42.4 / 30.0	33.1 / 17.7	64.2 / 44.5	52.6 / 35.9	50.0 / 32.7	27.7 / 14.1
MF <sup>10</sup> <sub>NEURAL</sub> -LEX	45.1 / 32.9	36.8 / 25.1	41.0 / 29.0	32.1 / 17.0	62.3 / 42.9	45.2 / 29.6	46.5 / 29.8	29.6 / 15.3

(a) Zero-shot DP scores. Source language: **English**.

	Seen Languages			Unseen Languages but Covered Scripts				New Script am
	hi	ur	Avg	bh	myv	wo	Avg	
	UAS / LAS	UAS / LAS	UAS / LAS	UAS / LAS	UAS / LAS	UAS / LAS	UAS / LAS	
mBERT	53.1 / 28.7	38.7 / 19.4	45.9 / 24.1	37.8 / 16.8	33.0 / 16.1	28.2 / 9.2	33.0 / 14.0	7.7 / 0.6
MAD-X 2.0	48.2 / 25.0	39.2 / 19.1	43.7 / 22.0	33.4 / 14.8	52.5 / 32.4	37.6 / 20.1	41.2 / 22.5	12.6 / 9.8
EL-RAND	46.2 / 22.7	34.1 / 15.7	40.1 / 19.2	30.0 / 10.0	49.6 / 25.8	22.7 / 6.9	34.1 / 14.2	27.3 / 11.0
MF <sup>1</sup> -RAND	48.0 / 25.5	36.0 / 17.7	42.0 / 21.6	29.9 / 12.2	54.5 / 32.0	34.6 / 17.3	39.7 / 20.5	33.8 / 13.7
MF <sup>10</sup> <sub>KMEANS</sub> -RAND	46.5 / 25.0	35.7 / 17.8	41.1 / 21.4	29.5 / 11.4	54.9 / 32.6	22.8 / 11.4	35.7 / 18.5	34.3 / 13.3
MF <sup>10</sup> <sub>NEURAL</sub> -RAND	44.3 / 22.9	31.7 / 15.3	38.0 / 19.1	32.2 / 13.6	53.6 / 31.7	29.9 / 15.6	38.6 / 20.3	32.9 / 11.7
EL-LEX	47.1 / 24.6	36.9 / 17.7	42.0 / 21.1	31.6 / 14.0	53.5 / 29.6	40.3 / 20.9	41.8 / 21.5	34.4 / 10.8
MF <sup>1</sup> -LEX	47.4 / 26.1	36.8 / 18.3	42.1 / 22.2	31.3 / 13.9	54.3 / 32.7	39.5 / 21.2	41.7 / 22.6	37.6 / 14.8
MF <sup>10</sup> <sub>KMEANS</sub> -LEX	47.6 / 25.8	38.4 / 18.9	43.0 / 22.4	33.5 / 14.3	54.5 / 32.7	40.5 / 22.5	42.8 / 23.2	28.4 / 11.2
MF <sup>10</sup> <sub>NEURAL</sub> -LEX	45.0 / 24.8	36.7 / 18.4	40.9 / 21.6	32.2 / 14.6	53.0 / 31.3	35.0 / 19.1	40.1 / 21.7	34.2 / 15.4

(b) Zero-shot DP scores. Source language: **Chinese**.

	Seen Languages			Unseen Languages but Covered Scripts				New Script am
	hi	ur	Avg	bh	myv	wo	Avg	
	UAS / LAS	UAS / LAS	UAS / LAS	UAS / LAS	UAS / LAS	UAS / LAS	UAS / LAS	
mBERT	57.0 / 40.7	43.7 / 29.8	50.3 / 35.2	45.9 / 30.6	27.7 / 9.3	18.9 / 2.3	30.9 / 14.0	10.0 / 0.9
MAD-X 2.0	55.5 / 38.3	45.4 / 30.4	50.5 / 34.3	42.6 / 27.2	37.9 / 19.7	22.2 / 6.1	34.2 / 17.7	12.3 / 9.7
EL-RAND	55.7 / 38.4	44.5 / 28.1	50.1 / 33.2	42.4 / 24.9	31.8 / 9.5	16.3 / 1.9	30.2 / 12.1	36.9 / 14.3
MF <sup>1</sup> -RAND	53.1 / 36.7	42.7 / 27.4	47.9 / 32.1	38.5 / 23.9	36.6 / 14.6	21.1 / 4.3	32.0 / 14.3	36.6 / 15.2
MF <sup>10</sup> <sub>KMEANS</sub> -RAND	54.6 / 38.0	44.6 / 28.7	49.6 / 33.3	41.3 / 25.9	36.6 / 15.5	14.1 / 2.3	30.7 / 14.6	33.1 / 14.0
MF <sup>10</sup> <sub>NEURAL</sub> -RAND	52.2 / 35.8	40.9 / 26.7	46.6 / 31.3	42.9 / 27.5	35.5 / 15.0	18.8 / 4.3	32.4 / 15.6	29.0 / 10.2
EL-LEX	52.9 / 36.3	45.1 / 29.1	49.0 / 32.7	40.6 / 27.1	34.9 / 13.3	22.6 / 5.3	32.7 / 15.2	37.6 / 11.6
MF <sup>1</sup> -LEX	53.1 / 36.8	44.2 / 28.6	48.7 / 32.7	41.5 / 27.0	36.1 / 14.8	23.9 / 5.6	33.8 / 15.8	31.2 / 12.5
MF <sup>10</sup> <sub>KMEANS</sub> -LEX	53.4 / 37.8	45.7 / 30.4	49.6 / 34.1	41.9 / 27.0	36.9 / 16.2	23.6 / 5.6	34.1 / 16.3	40.9 / 16.5
MF <sup>10</sup> <sub>NEURAL</sub> -LEX	53.5 / 37.9	44.5 / 29.2	49.0 / 33.5	41.3 / 27.0	35.6 / 16.7	20.4 / 4.8	32.4 / 16.1	32.6 / 14.5

(c) Zero-shot DP scores. Source language: **Japanese**.

	Seen Languages			Unseen Languages but Covered Scripts				New Script am
	hi	ur	Avg	bh	myv	wo	Avg	
	UAS / LAS	UAS / LAS	UAS / LAS	UAS / LAS	UAS / LAS	UAS / LAS	UAS / LAS	
mBERT	25.1 / 14.2	20.9 / 8.9	23.0 / 11.6	18.9 / 8.7	24.5 / 12.3	29.9 / 9.1	24.4 / 10.0	7.3 / 1.1
MAD-X 2.0	19.8 / 9.8	17.9 / 7.0	18.8 / 8.4	18.8 / 6.7	39.7 / 23.7	31.2 / 13.6	29.9 / 14.6	19.4 / 1.5
EL-RAND	19.7 / 10.0	15.0 / 6.0	17.4 / 8.0	17.5 / 6.6	42.2 / 25.5	25.2 / 8.9	28.3 / 13.7	30.0 / 6.7
MF <sup>1</sup> -RAND	20.3 / 9.4	18.4 / 7.3	19.4 / 8.4	19.7 / 5.3	36.7 / 20.7	33.1 / 13.3	29.8 / 13.1	34.0 / 8.9
MF <sup>10</sup> <sub>KMEANS</sub> -RAND	25.1 / 11.9	20.9 / 7.9	23.0 / 9.9	21.0 / 5.5	40.0 / 23.8	25.1 / 10.0	28.7 / 13.1	32.7 / 7.6
MF <sup>10</sup> <sub>NEURAL</sub> -RAND	23.8 / 11.4	15.3 / 6.3	19.5 / 8.9	19.1 / 6.1	40.3 / 22.6	31.2 / 12.7	30.2 / 13.8	0.0 / 5.8
EL-LEX	20.5 / 10.6	15.8 / 6.8	18.1 / 8.7	19.0 / 6.7	43.3 / 26.9	37.2 / 16.6	33.2 / 16.8	31.4 / 8.3
MF <sup>1</sup> -LEX	22.0 / 10.5	16.9 / 7.1	19.5 / 8.8	18.9 / 7.3	42.8 / 25.6	36.9 / 16.1	32.9 / 16.3	35.4 / 8.0
MF <sup>10</sup> <sub>KMEANS</sub> -LEX	25.2 / 12.8	21.7 / 8.8	23.4 / 10.8	15.1 / 5.6	31.6 / 17.0	38.0 / 17.5	28.2 / 13.4	28.7 / 6.6
MF <sup>10</sup> <sub>NEURAL</sub> -LEX	25.3 / 12.5	21.0 / 8.3	23.2 / 10.4	18.9 / 6.8	39.9 / 23.5	34.5 / 15.9	31.1 / 15.4	32.4 / 7.7

(d) Zero-shot DP scores. Source language: **Arabic**.

Table 11: Mean UAS and LAS test results for UD averaged over 5 runs transferring from the high-resource source languages (a) **English**, (b) **Chinese**, (c) **Japanses**, and (d) **Arabic**. The top group (first two rows) includes models which leverage the original tokenizer which is not specialized for the target language. The second group (last four rows) include models with new tokenizers. Here we separate models with randomly initialized embeddings (\*-RANDINIT) from models with lexical initialization (\*-LEXINIT) by the dashed line.

Language	Language code	Script	Language	Language code	Script
Afrikaans	af	Latin	Latvian	lv	Latin
Albanian	sq	Latin	Lithuanian	lt	Latin
Arabic	ar	Arabic	Lombard	lmo	Latin
Aragonese	an	Latin	Low-saxon	nds	Latin
Armenian	hy	Armenian	Luxembourgish	lb	Latin
Asturian	ast	Latin	Macedonian	mk	Cyrillic
Azerbaijani	az	Latin	Malagasy	mg	Latin
Bashkir	ba	Cyrillic	Malay	ms	Latin
Basque	eu	Latin	Malayalam	ml	Malayalam
Bavarian	bar	Latin	Marathi	mr	Devanagari
Belarusian	be	Cyrillic	Minangkabau	min	Latin
Bengali	bn	Bengali	Mongolian	mn	Cyrillic
Bishnupriya-manipuri	bpy	Bengali	Nepali	ne	Devanagari
Bosnian	bs	Latin	Newar	new	Devanagari
Breton	br	Latin	Norwegian-bokmal	nb	Latin
Bulgarian	bg	Cyrillic	Norwegian-nynorsk	nn	Latin
Burmese	my	Myanmar	Occitan	oc	Latin
Catalan	ca	Latin	Persian	fa	Arabic
Cebuano	ceb	Latin	Piedmontese	pms	Latin
Chechen	ce	Cyrillic	Polish	pl	Latin
Chinese-simplified	zh-Hans	Chinese	Portuguese	pt	Latin
Chinese-traditional	zh-Hant	Chinese	Punjabi	pa	Gurmukhi
Chuvash	cv	Cyrillic	Romanian	ro	Latin
Croatian	hr	Latin	Russian	ru	Cyrillic
Czech	cs	Latin	Scots	sco	Latin
Danish	da	Latin	Serbian	sr	Cyrillic
Dutch	nl	Latin	Serbo-croatian	hbs	Latin
English	en	Latin	Sicilian	scn	Latin
Estonian	et	Latin	Slovak	sk	Latin
Finnish	fi	Latin	Slovenian	sl	Latin
French	fr	Latin	South-azerbaijani	azb	Arabic
Galician	gl	Latin	Spanish	es	Latin
Georgian	ka	Georgian	Sundanese	su	Latin
German	de	Latin	Swahili	sw	Latin
Greek	el	Greek	Swedish	sv	Latin
Gujarati	gu	Gujarati	Tagalog	tl	Latin
Haitian	ht	Latin	Tajik	tg	Cyrillic
Hebrew	he	Hebrew	Tamil	ta	Tamil
Hindi	hi	Devanagari	Tatar	tt	Cyrillic
Hungarian	hu	Latin	Telugu	te	Telugu
Icelandic	is	Latin	Thai	th	Thai
Ido	io	Latin	Turkish	tr	Latin
Indonesian	id	Latin	Ukrainian	uk	Cyrillic
Irish	ga	Latin	Urdu	ur	Arabic
Italian	it	Latin	Uzbek	uz	Latin
Japanese	ja	Japanese	Vietnamese	vi	Latin
Javanese	jav	Latin	Volapuk	vo	Latin
Kannada	kn	Kannada	Waray-waray	war	Latin
Kazakh	kk	Cyrillic	Welsh	cy	Latin
Kirghiz	ky	Cyrillic	West	fy	Latin
Korean	ko	Korean	Western-punjabi	lah	Arabic
Latin	la	Latin	Yoruba	yo	Latin

Table 12: List of languages used in the pretraining. Taken from Chung et al. (2020)